Crowdfunding and Blockchain Economics

DOCTORAL THESIS

Author: Michele FABI

Supervisor: Dr. Matthew B. ELLMAN

UAB

Universitat Autònoma de Barcelona

MAY 2021

Page intentionally left blank.

To my beloved wife Eréndira, my parents Antonella and Stefano and all members of my beautiful family (even the smallest ones).

Acknowledgments

To begin with, I owe a special thanks to my advisor Matthew Ellman for guiding me with care and dedication throughout the development of this doctoral thesis and for having been an extraordinary mentor over the long journey of my doctorate. I also thank Matthew for having been a point of reference and having stood by my side even during difficult times of the doctoral path. I hold a profound and sincere esteem for him as an intellectual and as a person.

I want to also thank all other people who have contributed to my personal growth as a doctoral student. I thank many academics and colleagues who contributed to making me passionate about research, besides providing valuable discussions on my papers. In particular, Jordi Caballé; Ramon Caminal; Xavier Cuadras; Ramon Faulí-Oller; Sjaak Hurkens; Joachim Jungherr; Carolina Manzano; Jordi Massó; Antonio Miralles; Hannes Mueller; Chara Papioti; Amedeo Piolatto; Hugo Rodríguez. I thank the audiences at UAB Microlab, IAE (Institut d'Anàlisi Econòmica) internal seminars, Ca' Foscari internal seminar, TSE digital seminar, SAEe (Alicante, 2019), ASSET (Athens, 2019), WIPE (Reus, 2019), EARIE (Barcelona, 2019), JEI (Madrid, 2019) who helped me gain expertise in presenting my research. I thank great TA's such as Andrii Parkhomenko, Lavinia Piemontese, Shahir Safi for giving me the confidence needed to face difficult exams.

I am very grateful to Luca Corazzini and Pietro Dindo for their kind hospitality at Ca' Foscari University of Venice between September 2019 and January 2020 and for the help provided in obtaining the international doctorate mention, besides the many valuable discussions that helped me develop my ideas.

My full gratitude also goes to the people that helped and supported me personally during my four years of PhD. I thank and dedicate this work to my beloved wife Eréndira for always supporting me with plenty of patience and love and to my parents Stefano and Antonella for providing me plenty of encouragement. I also thank the amazing friends that made these difficult four years of the doctorate more colorful and pleasant. I thank my PhD colleagues Bruno Conte; Alfredo Contreras; Sven Hanold; Gauvrav Metha; Tommaso Santini; Chaoran Sun; Sarah Zoi for being part of the nice experience of growing together. It would not have been the same without you.

My sincere thanks go to the faculty and staff at IDEA and Barcelona GSE for creating a nice, cohesive, and stimulating environment, in particular to Jordi Caballé for giving

me the opportunity to join IDEA in the first place in 2015; Inés Macho for her dedication and support during the job-market training; David Pérez and Maite Cabeza for taking the helm of IDEA in difficult COVID times; Francesc Obiols for admitting me in IDEA as PHD student and helping me obtaining needed grants; Àngels López and Mercè Vicente for their huge help with administrative processes since the time I was a master student.

Page intentionally left blank.

Contents

Page intentionally left blank.

List of Figures

Page intentionally left blank.

Preface

This doctoral dissertation is a collection of three essays on the economics of emerging digital technologies; in particular, crowdfunding and blockchain. These practical technologies are increasingly used and have entered the public debate because of new regulatory questions. Although my research is purely theoretical, it already provides preliminary answers to vital questions about sustainability and welfare, so it has clear practical implications.

Chapter 1, entitled "The Cryptocurrency Mining Dilemma", studies how transaction fees and seigniorage, equivalent to inflation in the cryptocurrency setting, affect cryptocurrency miner incentives and trade congestion. It is centred on the analysis of cryptocurrency miners who are the record-keepers of the blockchain: a digital ledger of all transactions made with a particular cryptocurrency. By exerting effort to record pending transactions on the blockchain, miners determine the speed at which buyers and sellers of consumption goods can settle their trade with cryptocurrency. Given that a major reason to settle payments in cryptocurrency rather than traditional methods is its speed of payment (or settlement time), it is essential that miners record transactions at a good pace to sustain a cryptocurrency economy.

Each cryptocurrency miner chooses how many new transactions to group into the next block of transactions to record on the blockchain and competes with the other miners in recording it first. The successful miner earns (i) transaction fees from traders seeking to record their transfers and (ii) the new cryptocurrency units or seigniorage created by block validation. Pro-rata transaction fees encourage miners to include more transactions per proposed block but larger blocks transmit more slowly, raising the risk of invalidation. Therefore, raising the seigniorage to fee ratio reduces the number of transactions recorded per block. This aspect has important policy implications because, when miners choose small blocks, congestion levels rise and eventually (consumption goods) trade breaks down, making the cryptocurrency unviable. So a cryptocurrency designer has to impose a minimum threshold on the ratio of seigniorage to fees. The chapter finally presents the optimal cryptocurrency design problem as the maximization of trade efficiency subject to miner participation constraints and these miner incentives.

Chapters 2 and 3 are part of a joint line of work with my academic advisor Matthew Ellman on dynamic aspects of all-or-nothing (AoN) crowdfunding. Chapter 2, "A Theory of Crowdfunding Dynamics", develops a dynamic game featuring crowdfunding bidders

endogenously inspecting campaigns to understand the time profiles of their bids and to characterize how design parameters affect success rates, welfare and profits. Inspection costs convert a homogenous Poisson process of bidder arrivals into a non-homogenous Poisson bidding process. A *decreasing pivotality effect* generates a downward sloping average funding profile for any *fixed* inspection cost; the general logic is that strategic complementarity falls as a deadline nears. A weakly *concave* distribution of inspection costs reinforces this negative slope via a *negative Jensen effect* of success prospect news. Conversely, *convexity* can generate an increasing profile via a *positive Jensen effect.* Since decreasing pivotality is often strongest in the early stages of a campaign while positive Jensen effects from activating reluctant bidders are strongest in the final stages, this predicts a U-shaped profile in line with crowdfunding data. Besides providing a positive analysis, the framework we develop characterizes success rates given any price, allowing us to solve optimal pricing and address other design questions.

From crowdfunding data, we observe that bid profiles exhibit sharp U-shapes with initial and final spikes. These arise when some bidders arrive before campaigns open, or bidders delay their decisions to just before the campaign expires, creating a final bidding spike if the campaign turns out to succeed. Chapter 3, "Crowdfunding with Endogenously Timed Moves", is a spin-off of Chapter 2 that adds bidders the possibility to postpone their action from the moment they arrive and revisit the campaign just before its deadline by using a "Remind-me" option provided by the crowdfunding platform. Delay options let bidders adapt to news about bidding by others who act on arrival, but free-riding generates excessive delay. Relative to simultaneity (SIM) and exogenous sequentiality (EXO), endogenous sequentiality (ENDO) from delay options *can* lower success rates and welfare because of free-riding. However, with cost heterogeneity, ENDO can instead raise welfare because low-cost types endogenously move earlier, allowing higher cost types to economize on their inspections. In addition, ENDO can raise success rates by activating higher-cost types. This last dissertation chapter fully characterizes equilibria when the threshold is two. In this special case, cost homogeneity generates a continuum of pure and mixed strategy equilibria, but the expected mass of potential delays is the same in all equilibria. Only the unique equilibrium in stationary strategies is robust to cost perturbations. ENDO, EXO and SIM yield identical bidder surplus but SIM generates higher success rates, so SIM is welfare optimal under homogeneity. With a generic threshold, introducing cost heterogeneity favours ENDO and EXO relative to SIM in terms of consumer surplus. SIM continues to maximize the campaign success rate unless the cost distribution has a thick upper tail. In terms of bidding profiles, allowing for pre-campaign arrivals and endogenous timing, the model predicts initial and final bidding spikes, creating a pronounced U-shape.

Chapter 1

The Cryptocurrency Mining Dilemma

1.1 Introduction

A cryptocurrency is a digital currency governed by a computer program and managed by a decentralized, free-entry network of record-keepers in charge of updating the *blockchain*: a ledger (registry) of the collective memory of the cryptocurrency, i.e. "who-owns-what". In most cryptocurrencies, e.g. Bitcoin and its successors, record-keepers are *cryptocurrency miners*; individuals that own a dedicated computer. Miners record a *block* (or set) of pending transactions from consumption good traders (buyers and sellers) by first performing a costly encryption that makes the blockchain secure and then transmitting the block to the miner network for its approval. In case two or more blocks record conflicting information, miners coordinate on the first transmitted block and discard the others. As a compensation for the cost of recording blocks, the protocol rewards miners with revenues from two policy instruments: per-transaction fees and seigniorage from the creation of new coins. Seigniorage is the only source of revenues from empty blocks that record no transactions and only create money. By increasing block size recording more transactions, fee revenues rise while seigniorage is fixed, but block transmission time reduces; hence, also the risk of block invalidation causing the loss of its associated revenues increases. This paper studies the effects of seigniorage and transaction fees on miner incentives and trade congestion. My model shows that the size of miner blocks depends positively on the ratio of fees to seigniorage. Hence, transaction fees cannot be fully substituted by seigniorage as they are essential to induce miners in creating large blocks so to maintain a high transaction speed. Also, sufficient seigniorage is needed for positive miner activity when no trader transaction is pending. So an optimal cryptocurrency design uses both policy instruments.

Large blocks imply a short payment settlement period encouraging consumption good

buyers to pay sellers using the cryptocurrency. Conversely, small blocks cause the queue
of pending cryptocurrency payments to become congested and sellers to charge higher
prices for the wait. In my model, excessive seigniorage causes miners to form only empty
blocks, without fulfilling their primary record-keeping task. As a result, transactions
have an infinite validation time, and cryptocurrency trade unravels. Congestion risks in
cryptocurrency trade are evident from Bitcoin data. Figure 1.1 shows that most Bitcoin
traders suffered long payment settlement delays in periods of high congestion, as the
median confirmation time is often below the spikes occurring in the average confirmation
time. In particular, during the first quarter of 2018, transactions took on average more
than two days before being included in a block. Taking into account that Bitcoin traders
wait for other five blocks to arrive before considering a payment safely received, payments
took on average more than a week during that period of time, and caused traders to switch
to alternative payment methods such as Paypal transfers.[1]
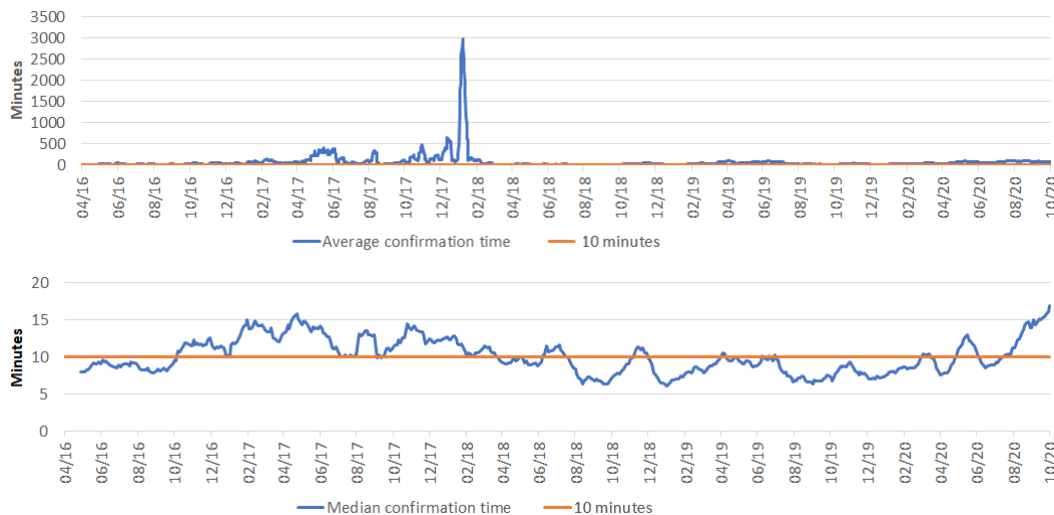
Figure 1.1: Average and median Bitcoin confirmation time

Day-level data origin: https://www.blockchain.com/
Data smoothing preformed via moving average with 10 lags.

The basic mining incentive problem can be tackled by an optimal cryptocurrency de-
sign that provides miners with the efficient combination of transaction fees and seigniorage
conditional on inducing sufficient miner entry for traders to consider the cryptocurrency
secure. In this case, the social planner (a team of expert software developers) elicits
"mining taxes" from consumption good traders in the form of inflation and transaction
fees to incentivize miners in providing public goods; in particular, security from their
participation, a pure public good, and block size, a common-pool good subject to con-

---

[1]On October 21 2020 Paypal announced to plan extending its service to allow trans-
fers of bitcoins and other cryptocurrencies across its accounts. For further information,
check https://newsroom.paypal-corp.com/2020-10-21-PayPal-Launches-New-Service-Enabling-Users-to-
Buy-Hold-and-Sell-Cryptocurrency

gestion. Chiu and Koeppl (2019) (CK hereafter) argue that a pure seigniorage design is optimal based on the observation that seigniorage is levied on aggregate value of the cryptocurrency while transaction fees apply only to the portion of value used for trade. The larger "tax base" on which seigniorage is charged can let the protocol provide a given level of mining revenues imposing lower mining costs on traders than the ones implied by including a fee component in the block reward design. However, CK's rationale neglects the detrimental effect of excessive inflation on block size that my paper explores. My analysis shows that an equilibrium in which cryptocurrency trade takes place is viable only if the protocol ensures that transaction fees are large enough relative to seigniorage so that miners are incentivized to fulfill their role of record-keepers.

My model is the first to jointly analyze congestion and miner incentives within a general equilibrium model of trade. Hence, my contribution to the Economics literature on blockchain and cryptocurrencies is threefold. First, I encompass a game-theoretic model of block mining within a state-of-the-art, continuous-time Lagos-Wright (LW) model, in this way endogenizing token demand and inflation. Then, I use the model to derive testable implications of the sensitivity of mining strategies and trade to shifts in policy variables; in particular, the mining fees-to-seigniorage ratio. Finally, I study optimal cryptocurrency design and contribute to the debate on the composition of the block reward in terms of seigniorage and transaction fees. The analysis is roughly consistent with the stylized facts on the co-movement between blocks' size and the fees-to-seigniorage ratio that emerges from the data (e.g. Fig. 1.2) and opens the way for further empirical investigation. Moreover, the main lessens form the model apply to the myriad of "altcoins" (alternative cryptocurrencies) that proliferated extending the open-source Bitcoin code.[2]

The rest of this paper is organized as follows: Section 1.2 provides a literature review. Section 1.3 models the cryptocurrency-trade economy and the blockchain, describing in detail the determinants of the block invalidation risk. Section 1.4 presents a game-theoretical model of mining and a partial equilibrium model of cryptocurrency trade. Cryptocurrency token prices are determined by the (general) monetary equilibrium of the economy presented in Section 1.5. Section 1.6 formulates the optimal cryptocurrency design problem and provides an intuitive suggestion on the composition of the block reward with a simple example. Concluding remarks and directions for further research are presented in Section 1.7. Omitted derivations and proofs are presented in Appendices 1.A

---

[2]An altcoin can be created by downloading and modifying the source code of the up-to-date version of Bitcoin Core from this link: `https://bitcoin.org/en/bitcoin-core/`.
Currently, the market capitalization of the more than 2800 different cryptocurrencies is around 200 million USD. Bitcoin is the market leader, accounting for 67% of the overall value with a capitalization of 133.43B. Ethereum and XRP (also known as Ripple) share the podium with market caps of 18.7B (9%) and 8.7B (4%) followed by the other top ten currencies (Tether, Bitcoin Cash, Bitcoin SV, Litecoin, EOs, Binance Coin, Tezos). A myriad of minor currencies constitute the remaining 7% of the total market value-the majority of each has a value share lower than 1%.
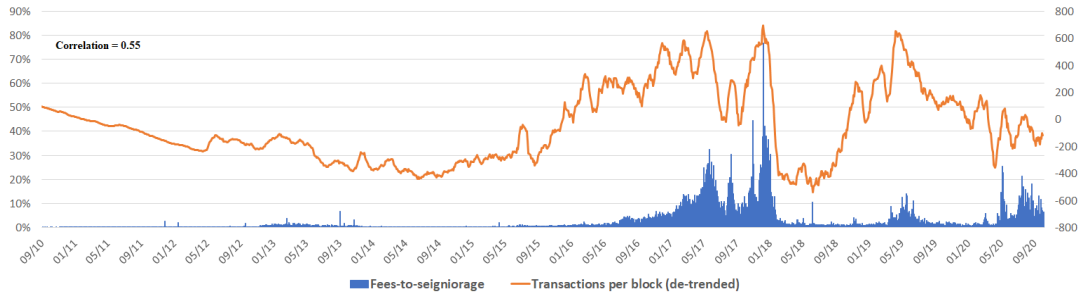
Figure 1.2: Fees-to-seigniorage ratio (in USD) and block transaction count for Bitcoin

Day-level data origin: https://www.blockchain.com/
Value "0" of the de-trended series corresponds to approximately 800 transactions.
Data smoothing performed via moving average using 3 lags for "Fees-to-seigniorage" and 10 lags for "Transactions per block (avg)"

and 1.B. Appendices 1.C and 1.D provide a summary of the main notation and technical terms.

1.2   Literature review

My paper contributes to the vibrant Economics literature on blockchain and cryptocurrencies. I develop a novel approach that encompasses endogenous block size, trade congestion, money demand and monetary policy within a general-equilibrium pure cryptocurrency economy based on the cutting-edge continuous-time adaption of the Lagos-Wright (LW) model proposed by Choi and Rocheteau (2020c) - CR hereafter. Lagos et al. (2014) provides a thorough review of the new monetarist economics literature of which CR is part.

Endogenous block size, congestion and inflation have already been studied separately and to some degree by other authors; yet, to my knowledge, the analysis of their interplay is new. In the closest paper to mine, Chiu and Koeppl (2019) develop a dynamic general equilibrium model of Bitcoin adoption where buyers (who are also miners) can engage in frauds by committing double-spending (DS) attacks. CK assume a fixed block size and show that the plausibility of DS attacks imposes a minimum requirement on miner participation and pins-down a compensation level that an optimal design has to guarantee.

CK advocate for a blockchain design based on a pure seigniorage block reward since it is more efficient in collecting revenues from traders than fees, but neglect congestion (see Section 1.1). CK also point out that fees can be reclaimed after a successful double-spending so that they provide a less effective DS deterrent than seigniorage. Nevertheless, in practice fees are small relative to the transaction value that justifies a DS attack and have only a minor effect on DS incentives. My story based on congestion and endogenous block size provides a rationale for adding a fee component to the block reward.

Houy (2016) along with the Computer Science paper Rizun (2015) is one of the earliest

contributions on endogenous block size with a model of the invalidation risk. He assumes continuous block size and also finds that block size is increasing in the fee-to-seigniorage ratio, in line with my findings. My paper augments the existing block size game by also studying miner activity and entry and also by endogenizing cryptocurrency trade. Earlier contributions miss modeling mempool dynamics resulting from the general equilibrium of the economy, which are an important link between mining and trade.

Huberman et al. (2019) and Easley et al. (2019) - HLM and EOB - offer an in-depth analysis of Bitcoin transaction fees. I focus instead on analyzing the optimal trade-off between fees and seigniorage, abstracting from the fee auction modeled by these authors and assuming homogeneous users and a fixed posted price as fee set by the protocol.[3]

Even though the models developed by EOB and HLM are suited to address the determination of fees, they treat trade as exogenous and assume linear impatience. This feature precludes a direct dynamic extension as standard models assume negative exponential discounting. On the other hand, CK endogenize trade and embed a game-theoretic model of mining within a LW monetary model in discrete time. This approach is suited for optimal cryptocurrency design but is not compatible with the continuous-time mining games studied by EOB and HLM; hence cannot be used to study congestion. My paper combines mining and monetary aspects within a continuous-time framework. As in EOB, I restrict the maximum block size to unity when investigating miners' optimal block size strategy. The novel mining aspect I consider is the trade-off between block reward and invalidation risk. HLM propose a modified Bitcoin protocol that adjusts block size and creation rate based on transaction demand. Block size is fixed by the protocol and miners do not take into account the risk of block invalidation.

Fernández-Villaverde and Sanches (2019) study a general equilibrium monetary model to examine market outcomes and welfare under (Hayekian) competition among private monies produced by profit-maximizing entrepreneurs. Similarly, Choi and Rocheteau (2020c,a) develop a monetary equilibrium model in which monies are created via costly (mining) technology. These authors find that stationary equilibria with steady-state inflation, in line with the approach I employ, exist. These equilibria are part of a multiplicity set featuring, for example, boom and burst dynamics and persistently declining purchase power. Schilling and Uhlig (2019) study a two-currency economy in which the US dollar and Bitcoin coexist. They determine a condition that rules out Bitcoin speculation and ensures that BTC price (in USD) follows a martingale. Specifically, Bitcoin speculation does not occur if agents are sufficiently impatient. Athey et al. (2016) develop a general equilibrium model of remittances to endogenize the Bitcoin price.

Biais et al. (2019), preceded by the Computer Science paper Kroll et al. (2013), use a stochastic game to investigate miners' fork resolution strategies. They demonstrate that LCR is an equilibrium mining strategy but coordination effects lead to multiplic-

---

[3]"Users" in HML and EOB are replaced by "traders" in my model.

ity of equilibria, some of them portraying permanent forking. Their paper includes an extension of the baseline model with delays in information transmissions, akin to the block transmission delays considered here, showing, in agreement with my analysis, that information delays can lead to temporary forks. The main modeling difference is that, in Biais et al. (2019), blocks are transmitted instantaneously to all miners except in one - and only one - transmission that can fail to reach a miner. In that case, the uninformed miner creates a fork as a result of the information asymmetry. In my model, blocks have different transmission times depending on their size, but completed block transmissions reach all miners at once. With my modeling approach, forking is caused by differences in transmission times. I first provide a detailed description of how LCR leads miners to discard (fork-out) blocks that have been transmitted slowly and then compute, in some cases explicitly, the probability of blocks becoming stale.

Prat and Walter (2018) estimate industry dynamics of Bitcoin mining contributing to the literature on irreversible investment and explaining price dynamics. They assume exogenous token demand and abstract from block reward design. Cong et al. (2019) provides an analysis of the industrial organization of mining pools.

Other papers, e.g. Abadi and Brunnermeier (2018); Leshno and Strack (2019); Saleh (2020); Rosu and Saleh (2019) are broadly related to mine and discuss general aspects of blockchain ecosystems and the Proof-of-Stake (PoS) protocol, the leading alternative to Proof-of-Work (PoW). In particular, Budish (2018) criticizes the blockchain technology highlighting potential vulnerabilities of the infrastructure to DS and other types of attack.

My paper also relates to the Computer Science literature on blockchains. Decker and Wattenhofer (2013) describe in detail the block propagation method used by Bitcoin miners and measure blocks propagation time on the Bitcoin blockchain. I refer to invalid blocks as "stale" according to the definition provided by Saad et al. (2019). Neudecker and Hartenstein (2019) study empirically temporary forks originated by block propagation delays. Carlsten et al. (2016) argue that a pure fee reward creates security breaches. In particular, random shifts in transaction fees caused by stochastic demand can cause mining revenues to fall below mining costs, thereby discouraging miner participation. This happens for example when no transaction is pending for validation so that miners make no revenues from mining. Rosenfeld (2014); Pinzón and Rocha (2016); Grunspan and Pérez-Marco (2018) compute the success probability of a double-spending attack refining the calculations reported in Nakamoto (2008).

1.3   Model

My model augments the novel continuous-time LW framework proposed by Choi and Rocheteau (2020c) with an explicit model of blockchain mining, featuring endogenous block space.

The economy is populated by B buyers, S sellers and M miners, indexed by b, s, m respectively. Traders' (buyers and sellers) participation is exogenous, while miner participation will be determined endogenously by free entry. I denote the time index $t \in \mathbb{R}_+$.

Two types of perishable and divisible goods are available in the economy. The first is a generic (numéraire) good denoted by $x \in \mathbb{R}$, with $x > 0$ if consumed and $x < 0$ if produced, that can be interpreted as a basic consumption good if positive or as labour if negative. The second is a special good whose consumption and production is denoted by $y \in \mathbb{R}$—it can be interpreted as a consumption good that is augmented with special features if purchased from an E-commerce website via cryptocurrencies. All agents enjoy consuming the generic good (and dislike producing it) according to the same one-to-one utility function, so the payoff they obtain by consuming or producing x is simply given by x. Preferences for the special good are instead asymmetric. Buyers cannot produce the special good but do enjoy consuming it according to the generalized logarithmic utility

$$u(y) = \ln\left(1 + \eta\, y\right), \quad \eta \in \mathbb{R}_+ \tag{1.1}$$

This functional form was first introduced in its more general form of generalized CRRA utility in Lagos and Wright (2005) and then employed by Chiu and Koeppl (2019) to normalize utility such that $u(0) = 0$ and to avoid a corner solution with no trade when even optimal consumption would yield a negative utility.[4] The taste parameter η is needed and has to be sufficiently large to solve this issue.

Sellers instead can produce the special good but do not enjoy consuming it, while miners are neither able to produce nor interested in consuming the special good.

All agents can also obtain storable and perfectly divisible tokens of a PoW cryptocurrency with no intrinsic consumption value, i.e. a fiat cryptocurrency. In Bitcoin, tokens are called "bitcoins" and "satoshis," (a bitcoin is worth 10^8 satoshis).[5] I let $a_{i,t}$, $i \in \{b, s, m\}$ denote the tokens held by agent i and $z_{i,t} \equiv a_{i,t}\phi_t$ their real value given price ϕ_t.

1.3.1 Centralized markets (CM's)

Only miners can produce cryptocurrency tokens. Nevertheless, all agents can obtain and dispose of tokens through centralized cryptocurrency markets (CM's). Specifically, two trading platforms, CM_1 and CM_2 (e.g. Coinbase and Binance), are continuously and simultaneously acting as market makers, allowing agents to exchange tokens for units of the generic good and vice-versa.

---

[4]The generalized CRRA used in Lagos and Wright (2005) is equivalent to $u(y) = (1 - \epsilon)^{-1}[(y + 1/\eta)^{(1-\epsilon)} - (1/\eta)^{(1-\epsilon)}]$, which converges to the utility function (1.1) for $\epsilon \to 1$.

[5]Fiat currency sometimes is interpreted as money backed by the government and issued by a central bank. I refer to these as *traditional* fiat currencies. Here, by "fiat currency" I refer to an asset that yields no dividends as opposed to a Lucas tree.

Besides being exchanges, the CM's allow traders to subscribe and opt for a token custody service. In this way, while retaining the tokens' ownership, traders let a CM store their tokens in compliance with the Know Your Costumer (KYC) and Anti Money Laundering (AML) regulations. Trading platforms perform the custodial service by providing traders with a *custodial wallet* that they can use to manage their funds. In this model, all traders opt for the custody service and operate only trough their custodial wallets. Miners instead keep their tokens "in their own hands" on a non-custodial wallet. To keep things simple, I assume all traders single-home and each CM interacts with half of the traders' population. Formally, denoting a CM's user base with sub-index, I am making the following assumption:

Assumption 1.1. $(B_1, S_1) = (B_2, S_2) = (B/2, S/2)$

Miners instead operates with both CM's.

Each CM performs two types of operations: internal transactions, among and with her subscribed traders, and external transactions, encompassing those among subscribed and unsubscribed traders and miner-platform transactions. Transactions of this last category take place after miners recorded them on the blockchain in a valid block, i.e. they are *on-chain* transactions. In this case, their settlement is not immediate, and record-keeping is in general not final because, once a block becomes stale, the transactions it records turn to be pending again. Conversely, internal transactions are performed using a double-layer solution that does not rely on the blockchain; i.e. they are *off-chain* transactions. To perform transactions of this category, first, a CM forms a centralized fund buying tokens via external transactions—at first buying from miners and later also from unsubscribed traders. Afterwards, she lets subscribed traders access a limited part of her fund ($a_{b,t}$ and $a_{s,t}$ to a generic buyer and seller) through their custodial wallets. In this way, internal transactions only update traders' wallet balance, but de facto do not move tokens across addresses. Thanks to this procedure, internal transactions achieve immediate and final settlement.[6]

All external transactions experience settlement delays, but I assume that miner-to-platform operations are immediate, as considering the effect of their settlement delays would only complicate the model without making it further insightful.

1.3.2   Decentralized market (DM)

Besides trading in the CM's, traders also meet sporadically among each other in a decentralized market (DM). Given the preference and technology asymmetry among buyers and sellers, DM meetings are the only occasion for buyers to purchase the special good

---

[6]Indeed, the Bitcoin blockchain explorer https://www.blockchain.com/explorer does not record transactions made through Coinbase wallets, as explained in Coinbase FAQ (frequently-asked-questions).

from sellers. I let $\alpha_{i,j}$, $i,j \in \{1,2\}$ denote the (Poisson) meeting rate among CM i buyers and CM j sellers.

Since the custodial wallets are incompatible across CM's, DM transactions among traders of different platforms occur *on-chain*; precisely, when a transaction takes place, the buyer requests his trade platform to send a number of tokens to the seller's address, which points to his trading platforms. In this model, a precise analysis of blockchain record-keeping is pointless unless traders with incompatible wallets meet with positive probability. So, to keep things tidy, I assume that DM meetings occur *only* among buyers and sellers with incompatible wallets, as stated by the following assumption.

Assumption 1.2. $\alpha_{1,2} = \alpha_{2,1} \equiv \alpha, \quad \alpha_{1,1} = \alpha_{2,2} = 0$

During DM meetings, I assume that credit is ruled out by market practices or anonymity - e.g. if special good is illegal, traders identify themselves using an encrypted alias.[7] Assuming that buyers cannot produce the generic good while trading with sellers, the only way they can obtain the special good is by exchanging it for their cryptocurrency tokens. This feature makes the cryptocurrency essential, in the sense that it widens the frontier of welfare-improving trade arrangements.

To reward miners for recording DM transactions on the blockchain, traders' wallets are programmed to charge a transaction fee for each transaction made. I assume that wallets charge a proportional transaction fee rate τ on each transaction. To be clear, if buyer b sends tokens of value $z_{b,t}$ to seller s, this latter receives only $z_{s,t} = z_{b,t}(1-\tau)$. The miner m that records the transaction between b and s in a valid block receives $z_{m,t} = z_{b,t}\tau$ real balances in reward for doing so.

In reality, buyers set the amount of fees attached to their transactions. Here, by considering τ as fixed, I give up some realism for the sake of simplicity. The model provides the simplest yet parsimonious setting to study the general equilibrium implications of a change in transaction fees.

Transaction fees are not the only source of rewards for miners and currently constitute only a minor component of it. The most important part of their reward comes from seigniorage in the form of new tokens that are generated by the protocol each time a miner forms a valid block. I postpone this aspect to Section 1.3.4 that describes the details of tokens' creation and the composition of the block reward.

Market clearing and symmetry

Each CM sets her token price $\phi_{i,t}$ to ensure market clearing at each date and for each realization of (tokens) demand and supply shocks. Formally, letting $(A_{i,t}^D, A_{i,t}^S)$ denote

---

[7]Even in these cases traders' identity can be retrieved indirectly by analyzing the blockchain tree. For example, The FBI was able to trace the identity of most people involved in illicit trade through the website Silk Road.
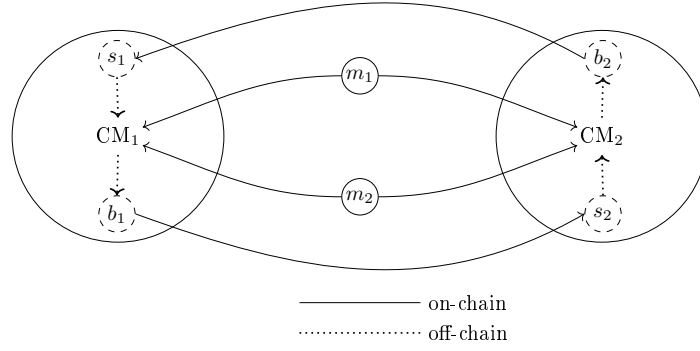
Figure 1.3: On-chain and off-chain transactions

CM i's token demand and supply, market clearing implies

$$\phi_{i,t} : A_{i,t}^D = A_{i,t}^S \equiv A_{i,t} \quad \forall\ i,t \tag{1.2}$$

The resulting aggregate real value of tokens is

$$Z_t \triangleq \phi_{1,t} A_{1,t} + \phi_{2,t} A_{2,t}$$

In the monetary (general) equilibrium characterized in Section 1.5, the token demand in each CM is constant - except at zero-measure time points - but the token supply is subject to shocks. Specifically, positive supply shocks occur each time a CM receives tokens from a miner-platform or DM transaction, while negative supply shocks occur every time a CM gives away her tokens in a DM transaction. Nonetheless, in expectation, the symmetry Assumptions 1.1 and 1.2 imply that DM shocks compensate each other across CM's. Therefore, CM's expected prices are equal and are determined by the (constant) tokens' demand and miners' token supply. Since traders' portfolio choice and mining strategies are based on the same expected price, we can without loss of generality drop the CM index when analyzing the partial equilibrium models of mining and cryptocurrency trade developed in Section 1.4, basing the analysis on a single price ϕ.[8]

The aggregate number of tokens in circulation varies over time, but the general equilibrium I will characterize in Section 1.5 has the property that aggregate real balances are constant in each CM.

Definition 1.1 (Stationarity property)**.** *Aggregate real balances are stationary at CM level if*

$$Z_{i,t} = Z_i \ \ for\ i = 1, 2 \tag{1.3}$$

---

[8]As in reality, realized prices can temporarily differ despite the symmetry assumptions, providing agents with arbitrage opportunities. However, arbitrage is not worthwhile if its expected gains are offset by the hassle cost of monitoring both CM's at the same time. I will take his aspect for granted hereafter.

and for all $t \in \mathbb{R}_+$ except for a set of dates with zero Lebesgue measure.

The equilibrium describes a situation in which tokens' (expected) price is on a steady inflation path and is not subject to speculative bubbles. In reality, BTC price volatility caused by speculation raises major concerns regarding the use of BTC as a means of payment (MoP). The two-platform structure of the model is suited for analyzing cryptocurrency speculation, but in this draft, I concentrate on the hypothetical scenario in which problems arising from speculation are resolved.

Summary

To conclude this section, I recap how cryptocurrency tokens and consumption goods circulate in the economy as illustrated by Fig. 1.4. In type-(i) exchanges, miners sell their block rewards to the CM's in exchange for generic goods. In the ones of type-(ii), buyers acquire tokens from CM's by selling units of the generic good and use their tokens to trade with sellers in the DM where type-(iii) exchanges take place: a share of the value transferred in DM meetings goes a sellers (iii.a); the remaining part goes back to a miner in the form of transaction fees (iii.b).[9] Type-(iv) exchanges occur as soon as DM transactions are recorded on a valid block, so that sellers cash-out their tokens by selling them in the CM in exchange for units of the generic good.
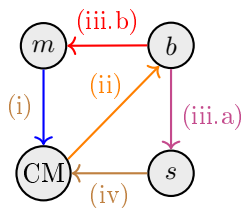
Figure 1.4: Circulation of a token

Understating the internal mechanisms ruling the blockchain is essential for studying the general equilibrium of the cryptocurrency economy. For this reason, the next section presents a detailed model of blockchain mining that complements the trade framework presented up to now.

1.3.3 The blockchain

The blockchain is a digital ledger that records all movements of cryptocurrency tokens directed towards each user address as well as modifications of the token creation policy and in the ledger's internal governance. The kind of blockchains studied in this paper are based on the PoW protocol that puts miners in charge of three fundamental tasks: record-keeping, consensus formation and security.

---

[9]Technically, transaction fees resulting from the difference between the number of tokens sent from an address and that directed to the other address or addresses.

Each miner stores a copy of the ledger on his mining node and records incoming transactions sent by traders. Transactions are recorded on a miner's blockchain copy in time-stamped batches called *blocks*. To form blocks, a miner is required to solve a computationally-intensive cryptographic puzzle through random guessing of its solution, an operation called *mining* that nowadays requires a dedicated ASICS hardware to be performed.[10] In reward for mining, miners receive seigniorage from the creation of new tokens and fees associated with each transaction they record.[11]

The costly PoW requirement exists for two main reasons. The first is to give value to the cryptocurrency by making its supply costly, as the consequences of removing the PoW mechanism on the value of the cryptocurrency are analogous to those of letting people print M0 and M1 money (basically coins and banknotes) with their home printer on the value of a traditional fiat currency. The second role is related to security, but I delay this aspect after having described the structure of a blockchain in the paragraphs below.

Due to the distributed nature of blockchain record-keeping, miners are naturally prone to record different transactions histories to their ledger copies. For example, if miners form blocks by recording pending transactions at random, their ledger copies can differ in the chronological order in which transactions are recorded. Nevertheless, the governance rules followed by miners have to ensure that miners reach an agreement on a common version of the ledger by communicating to each other the blocks they recorded and pending transactions they store, allowing only for temporary inconsistencies among their ledger copies.

In the next subsection, I will provide a brief description of the structure of a blockchain and the miner governance rule adopted by Bitcoin and most PoW cryptocurrencies. The following short description suffices to follow the mining model developed in Section 1.4.1. Other minor mining technicalities are presented in Appendix 1.D.

The structure of a blockchain (in brief)

A blockchain is formed by connected blocks of recorded transactions and auxiliary information. Each blockchain is initiated by a genesis block that is progressively extended by its successors establishing a chronological order. Asynchronous and decentralized communication among miners can lead them to extend a single block by two or more direct (chronological) successors. In this case, the blockchain *forks* (bifurcates) in two or more chains (branches), each providing a different version of the ledger up to a common point

---

[10]ASICS is an acronym for Application-Specific Integrated Circuit System. ASICS mining nodes were anticipated first by computational processing units (CPU's), and then by graphical processing units (GPU's).

[11]An alternative blockchain protocol is Proof-of-Stake (PoS) which replaces miners with "validators," who are required to form a token escrow fund and obtain the right to record blocks if extracted by a lottery that selects them based on their relative contribution to total token escrow.

of agreement. Due to possible forking, the correct model to keep track of the blockchain's ramification is that of a directed tree graph, usually referred to as the *block tree*, described rigorously by Biais et al. (2019). Identifying a block's precise location within the block tree requires modeling techniques that are not necessary to develop the analysis that follows. Hence, I identify each block only by the identity of the miner m who forged it and a block height $h \in \{0, 1, \dots, H_t\}$ that counts its block distance from the genesis block ($h = 0$). The variable H_t indicates the blockchain height (i.e. the height of the block head of the longest chain). I omit the time index from H_t when clear form the context or referring to a generic value H of the blockchain height. I also refer to a chain's height as the height of its block head.

Each chain (branch) of the block tree is initiated either by the genesis block or by a fork and portrays a different history of the ledger. For this reason, the presence of multiple active chains create ambiguity on the state of the ledger, e.g. on the balance associated with each cryptocurrency address, and hinders users' trust in the blockchain if persistent. To avoid a state of permanent ambiguity, the governance rules followed by miners have to guarantee that they "form a consensus" over a unique chain to extend, allowing for simultaneous active chains only temporarily.

Bitcoin and most (if not all) PoW blockchains follow the chain selection criterion prescribed by Nakamoto (2008); namely, the Longest-Chain-Rule (LCR). Essentially, LCR prescribes miners to consider as valid the *perceived* longest chain by electing its block head (terminal block) as reference (predecessor) block for the new block they will form. If a single block has height H, miners following LCR form their new block on top of it. On the other hand, if the blockchain terminates with a fork (multiple blocks have height H), miners work on the terminal block they became first aware of; that is, on the block with the shortest *publication time* (the stopping time at which it is recorded on the ledger and transmitted to the miners' network). Put differently, miners following LCR behave as if they were to choose their favourite chain according to a lexicographic preference, with chain height as primary criterion and publication time of a chain's block head as subordinate criterion.

Under LCR, a miner m^* extends the blockchain at height H establishing the new consensus chain (at height $H + 1$) by publishing a block with the shortest publication time T_{H+1,m^*} among the ones of all the miners $m \in \mathcal{M}$ that are simultaneously competing to extend the longest chain. Formally, LCR determines the reference block as follows:

Criterion 1.1 (LCR)**.** *Under the Longest-Chain-Rule (LCR), a miner m^* establishes the consensus chain at height H if*

$$T_{H+1,m^*} = T_{H+1}^* \triangleq \min \left\{ T_{H+1,m} \right\}_{m \in \mathcal{M}} \tag{1.4}$$

From the perspective of a miner following LCR, all blocks outside the consensus chain
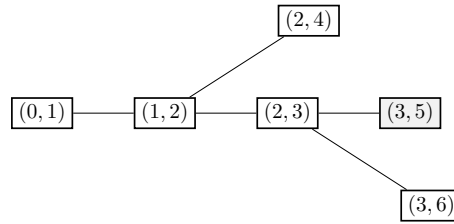
Figure 1.5: Longest Chain Rule (LCR)

A branched block tree. Blocks are identified with height (left-index) and the rank of their publication time (right-index). Among the three chains that form the block tree, LCR selects the longest chain terminating with block $(3,5)$. Notice that the chain ending with block $(3,6)$ is one of the longest but was published later than the consensus chain, whereas the chain ending with block $(2,4)$ was published previously than the consensus chain but is shorter.

are *stale* and do not contribute to the ledger's recorded history of transactions.

LCR is a natural prescription under transaction homogeneity. If we depart from this benchmark case, there are reasons to believe that rational miners can follow a different behaviour. For example, with heterogeneous transaction fees, Carlsten et al. (2016) shows mining equilibria such that miners extend the branch whose terminal block provides them with the largest amount of transaction fees. Also, Biais et al. (2019) shows that LCR is a plausible mining prediction, but also that other equilibria can arise due to coordination motives, e.g. prescribing miners to extend the last longest branch they become aware of (with the longest update time) rather than the first. The same coordination motives that sustain LCR can also lead to an equilibrium with permanent forking.

Miners also deviate from LCR by intentionally forking the blockchain when committing a fraud, e.g. altering the history of recorded transactions - as in the case of history reversals, including DS attacks, described later in Section 1.3.4 - or when using in selfish mining strategies analyzed by Eyal and Sirer (2014).

1.3.4 Mining

To form a new block, a miner has to (i) select a subset of pending transactions from his mem-pool (ii) choose a reference predecessor and find a PoW for his block (iii) communicate his block successfully to the other miners.

PoW and block rate

Each miner m can generate PoW solutions at a Poisson rate μ_m determined by the hash-power (production of PoW solutions) produced by his mining rig and the PoW difficulty set by the protocol. Conversely, the (average) block creation rate

$$\mu \triangleq \sum_m \mu_m$$

is fixed by the protocol, which scales-up the PoW difficulty based on the aggregate hash power. In Bitcoin, a difficulty adjustment is applied every two weeks to restore an average block creation time of 10 minutes based on an estimate of miners' aggregate computing power (see Fig. 1.6).
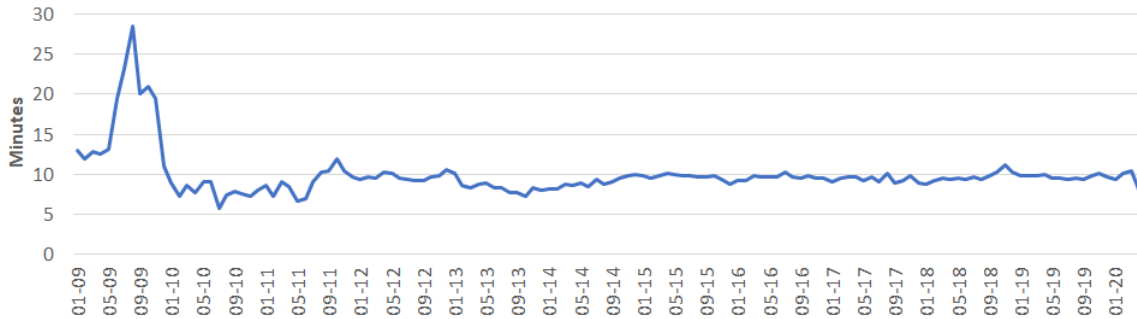
Figure 1.6: Block mining time (average)

Month-level data from https://data.bitcoinity.org/bitcoin/block_time/all?f=m10&t=l

Biais et al. (2019) provides a detailed description of how the difficulty adjustment is implemented. Under hash rate homogeneity, a generic miner's PoW rate as a function of the number of other active miners M is described by the formula[12]

$$\mu_m = \frac{\mu}{M} \quad \forall m \tag{1.5}$$

From the $1/M$ term in the above expression we can see that each miner creates a negative difficulty externality on the other miners.

For a steady-state distribution of miners' mem-pools to exist, the rate at which transactions are processed has to be higher than transaction request rate $\lambda \equiv \alpha B$. The following condition ensures that miners' mem-pools do not explode when the processing capacity they offer is maximal.

Assumption 1.3. $\mu > \alpha B$

Transmission delays

After finding a valid PoW, a miner can collect a block reward for his block only if it does not become *stale* in an abandoned branch of the blockchain.[13] In my model, delays in miners' communication are the only source of stale blocks.

To keep thing simple, I assume that users communicate their transaction requests instantly to all miners and each miner observes the completion time of his block transmissions, which reach all other miners simultaneously. The only communication friction is

---

[12]In what follows mining activity is constant over time so that also PoW rates are constant. With time-varying aggregate mining power also hash rates would fluctuate.

[13]Some authors refer to blocks outside the consensus chain as "orphan blocks". I follow Saad et al. (2019) and consider as orphan the blocks that extend an invalid predecessor.

provoked by transmission delays that depend on the size of blocks. In this model, blocks are either empty or filled with one transaction only. For now, to facilitate exposition, I denote the size of a miner's block by $k_m \in \{0, 1\} \cup \{\text{OFF}\}$, where I use the convention that $k_m = \text{OFF}$ if the miner does not produce any block (his mining rid is shut down). When presenting the game-theoretic model of mining in Section 1.4.1 I will let miners strategize on the size of their blocks based on number Q_t of pending transactions in their mempools, so that they will set $k_m(Q_t)$ for $Q_t \in \mathbb{N}_0$.

Empty blocks are transmitted immediately, while filled blocks are transmitted after an exponentially distributed transmission lag with average $1/\theta$. I further assume for tractability that miners update their ledger copies only after the blocks they create are published (mined and transmitted). Under these assumptions and LCR (described in Section 1.3.3), all miners work on extending the unique consensus chain.

Apart from the case in which $k_m = \text{OFF}$, the time it takes a miner m to forge and transmit a new block that *extends* the consensus chain at height H is the sum of the PoW solution time $\gamma_{H+1,m}$ and transmission lag $\epsilon_{H+1,m}$ of his new block. Thus, its the publication time results from

$$T_{H+1,m} = T_H^* + \gamma_{H+1,m} + \epsilon_{H+1,m}$$

where T_H^* is the publication time of the reference block for height H. The solution time $\gamma_{H+1,m}$ is exponentially distributed with rate μ_m in accordance with Eq. (1.5). The transmission lag $\epsilon_{H+1,m}$ of miner m's new block depends on its size. If $k_m = 0$, its transmission is instantaneous ($\epsilon_{H+1,m} = 0$); if $k_m = 1$ its transmission lag is exponentially distributed with rate θ. The density of the total solution and transmission time when a miner fills his new block with probability σ_m is the given by the following mixture distribution:

$$f_{\gamma_{H+1,m}+\epsilon_{H+1,m}}(t) = \sigma_m \left[\frac{\theta \mu_m}{\theta - \mu_m} \left(e^{-t\mu_m} - e^{-t\theta} \right) \right] + (1 - \sigma_m)\mu_m e^{-t\mu_m} \qquad (1.6)$$

Notice that, for $\sigma_m = 0$, density (1.6) becomes a simple exponential density, while in the opposite case of $\sigma_m = 1$, it becomes the density of the sum of two exponential random variables with different rates (a two-parameter hypo-exponential distribution).

Under LCR, all miners will use as next reference block the one extending the longest chain with the shortest publication time according to relationship (1.4). Hence, the probability that a miner m^* forms next reference block satisfies

$$\mathbb{P}\left(T_{H+1}^* = T_{H+1,m^*}\right) = \mathbb{P}\left(\gamma_{H+1,m^*} + \epsilon_{H+1,m^*} < \min_{m'}\left\{\gamma_{H+1,m'} + \epsilon_{H+1,m'}\right\}\right) \qquad (1.7)$$

If successful, miner m^* establishes the new longest chain and causes all other blocks in transmission to become stale, each forming a separate abandoned chain. Formula (1.7) is

Figure 1.7: Three miners updating a blockchain

Blocks are indexed $(h - h_0, m)$ for $m \in \{1, 2, 3\}$ and $h - h_0 \in \{0, 1, 2, 3, 4, 5, 6\}$, with $h_0 > 0$. Forks occur at $h - h_0 \in \{3, 5\}$ and are resolved according to LCR that selects as reference block among the head blocks of competing (longest) chains the one with the shortest update time (in grey).

employed in Section 1.4.1 to determine a miner's estimate of the probability of successful mining of a block given a belief on the size of the other blocks in contemporaneous transmission. Fig. 1.7 shows an illustrative blockchain ramification caused by three miners competing to establish the consensus chain given the characteristics of block transmission stylized in this section.

In current cryptocurrency blockchains, information transmission is quick and encounters little geographical impediments.[14] The resulting transmission delays are of a lower order of magnitude than block creation times. To be precise, according to the estimates for Bitcoin reported in Decker and Wattenhofer (2013), $\mu \approx 1/600$ (one block creation every 10 minutes) and $\theta \approx 1/10$ (one completed transmission every 10 seconds). These numbers are in line with Fig. 1.6 on block creation and Fig. 1.8 on block propagation. Therefore, the effect of transmission lags on the block creation process can be safely ignored, as well as the probability that a miner finds a PoW before completing a block transmission.[15] Based on the previous observation, I approximate the distribution of inter-update time of reference blocks using an exponential distribution with parameter μ:

$$f_{T^*_{H+1} - T^*_H}(t) \approx \tilde{f}_{T^*_{H+1} - T^*_H}(t) = \mu e^{-\mu t}, \quad t \geq 0 \tag{1.8}$$

This approximation can lead to slightly different probabilities than the ones produced by the actual distribution of the inter-update time between reference predecessor blocks but does not affect any qualitative result of this paper. Moreover, it allows to solve explicitly and neatly all agents' value functions, which otherwise would be too complicated

---

[14]The Canadian blockchain technology company Blockstream started recently to broadcast the Bitcoin blockchain via satellite to facilitate the transmission and reception of Bitcoin data from areas with scarce or inaccessible internet services. More information is available at the following link: https://blockstream.com/satellite/

[15]A block transmission is considered completed when it reaches a certain majority of miners. Fig. 1.8 shows that a block reaches 90% of Bitcoin miners' population in about 10 minutes.

and convoluted to be handled.[16]

Figure 1.8: Block propagation time (% of total miners reached)*

*Average time until a block is announced by a given percentage of total Bitcoin miners.

Day-level data origin: https://dsn.tm.kit.edu/bitcoin/#propagation
Data smoothing preformed via moving average with 15 lags.
All observations are averages of block propagation times recorded within a time span of approximaly one hour recorded between 12AM and 2PM.

In reward for successful mining, miners earn a block reward $R(k_m)$ depending on the size of their blocks, with $R(1) \geq R(0) \geq R(\text{OFF}) = 0$. By choosing to record a filled rather than empty block, a miner faces a higher risk that his block becomes stale; but, if the block transmits successfully, he earns a higher reward from transaction fees.

Block reward

Each time a miner successfully records a valid block, he earns a block reward composed of seigniorage from newly created tokens and transaction fees. To award a miner with the latter, the protocol transfers new tokens to his address with a *coinbase* transaction. Metaphorically speaking, miners seek the "digital gold" contained in each valid block. Hereafter, each (valid) block contains an amount ΔA_t of new tokens, set in fixed proportion to the total token supply.

Since cryptocurrency trade occurs at sufficiently high frequency, a law of large number (LLN) ensures that trade supply shocks are neutralized at CM aggregate level. In this way, price shocks are only caused by miners selling their coinbase rewards. Throughout this paper, I let the protocol set a steady inflation rate in both CM's. Specifically, if a miner sells a coinbase transaction to the CM's at time t_+, the resulting price shock in a generic CM has a magnitude of $\phi_{t_+}/\phi_t = 1 - \pi$, where $\pi \in [0, 1]$ denotes a seigniorage rate. The protocol implements the desired seigniorage rate π by setting

$$\Delta A_t = A_t \frac{\pi}{1 - \pi}$$

Under steady inflation, the value of miners' seigniorage reward is the product of the seigniorage rate π and the stock of aggregate real balances Z_t.

---

[16]The exact distribution of μ is the one of the minimum of M random variables with density (1.6).

$$\Delta A_t \phi_t = Z_t \pi \tag{1.9}$$

The block reward $R(k_m)$ combines seigniorage with transaction fees of rate τ earned on the k transactions included in a block. Given that, in equilibrium, all B buyers carry the same amount of real balances $z_t = z$, so that $Z = zB$, the total block reward amounts to

$$R(k_m) = z\left(\pi B + \tau k_m\right), \; k_m \in \{0, 1\}, \qquad R(\text{OFF}) = 0 \tag{1.10}$$

The token creation rule presented before establishes a constant inflation rate. Bitcoin implements a staggered decreasing inflation rule such the tokens produced in each valid block are halved each time the total token production hits one of a set of predetermined target values. These events are called "Bitcoin halvings" and are programmed to impose a total supply cap of 21 million bitcoins, to be reached approximately in the year 2140.[17] After then, Bitcoin will feature block rewards made only by transaction fees and a negative inflation rate driven by tokens going out of circulation by getting lost or ending up in abandoned accounts.

The model applies to Bitcoin in a scenario where the next halving is far from agents' temporal planning horizon. It is also valid for the many cryptocurrencies based on a steady token creation rule, such as the current PoW implementation of Ethereum. It does not apply to cryptocurrencies that implement a negative inflation rate by "burning tokens," sending a part of them to an irretrievable address each time a valid block is recorded.

Security

Suppose that the most extended blockchain branch has height H_t and a malicious miner deviates from LCR and attempts to re-write the information stored in a block at height, $h < H_t$. Since block identifiers are chained recursively, to succeed in doing so by time t' the hacker has to create a secret chain that re-writes all blocks at height $h' \in \{h, h + 1, \ldots, H_t, \ldots, H_{t'}\}$, where $H_{t'} \geq H_t$ is the blockchain height reached during the attack, solving the PoW cumulatively, and release his branch once it surpasses the honest one. In this way, the attacker engages in a mining race with the honest miners, as he needs to keep up with the additional blocks that extend the blockchain. The likelihood that the attacker succeeds is decreasing with the participation of honest miners because it makes honest (longest) chain to grow faster. Hereafter, as in Easley et al. (2019), I assume that cryptographic attacks are ruled out if at least $\underline{M}$ (honest) miners are active. This level of mining activity is also required for traders to trust in using the cryptocurrency.

Now that all the fundamental ingredients of the cryptocurrency-based economy have been introduced, I am ready to move to the next section, in which I analyze the partial

---

[17]The last Bitcoin halving took place on May 11th, 2020.

equilibrium models of mining and trade.

1.4 Mining and trade equilibria

In this section, I present the partial equilibrium models of mining and trade. Mining involves strategic interactions that require game-theory modeling. Essential elements for the design results discussed in Section 1.6 are the probability at which a given miner works on recording a block (1.16) and miner participation under free-entry (1.25). I also determine the necessary conditions for any mining activity to occur (1.19) and a sufficient condition for each miner to be constantly mining (1.20).

Cryptocurrency trade can be tackled using dynamic programming (DP) and axiomatic bargaining. I use these techniques to derive the optimal amount of special good produced by sellers (1.35) and buyers' token demand (1.36).

1.4.1 The mining game

Section 1.3.3 described the technicalities and purposes of mining and the PoW protocol. Here I develop a stochastic game to analyze miners' block size choice, use of computational power and participation to the mining network. I present its elements hereafter:

Players: M miners compete for updating the consensus chain of a blockchain recording cryptocurrency transactions.

States: Miners engage in mining rounds (or tournaments) with features varying according to a (continuous-time) Markov chain. Precisely, each round is described by a two-dimensional state, given by a number of pending transactions and a blockchain height, $(Q, H) \in \mathbb{N}_0 \times \{H_0, H_0 + 1, \dots\}, H_0 > 0$.

Information: A generic miner m observes his copy of the ledger continuously as well as the PoW solution time $\gamma_{h,m}$ and the transmission lag $\epsilon_{h,m}$ of all the blocks he mined and transmitted. By combining this information, the miner learns the update time of his blocks $T_{h,m}, \forall h \leq H + 1$.

The miner also observes the update times and the information contained in the blocks transmitted by other miners $T_{h,m'} \ \forall h \leq H + 1, m' \neq m$, except for the other miners' identities. Thus, since he observes all update times, he also knows the publication time of all reference blocks T_h^*, again for $h \leq H + 1$.

Finally, the miner observes his mem-pool size $Q_{m,t}$ at each point in time. Since users-to-miners communication is simultaneous and immediate and miners update their mem-pools only when either receiving transactions or recording them in finalized valid blocks, miners' mem-pools are always synchronized, recording the same number of pending transactions $Q_{m,t} = Q_t \ \forall m$ that I denote simply as Q when the temporal dimension is implicit.

Actions: Each miner m chooses the size of the block $k_m(Q_t) \in \{0, 1\} \cup \{\text{OFF}\}$ to mine on top of the longest chain and transmit to the miners' network afterward. I adopt the convention that $k_m(Q_t) = \text{OFF}$ if miner m is idle (his mining rig is switched off) when the mem-pool has size Q_t. I let $Q_{T_m^* + \gamma_m}$ denotes the state of the mem-pool by the time the miner finds a PoW for a block extending the longest chain.

Payoffs: The lifetime utility of a miner is the present value of the sum of the block rewards he earns in each mining round net of the upfront investment cost for his mining node F and a flow mining cost ψ accounting for energy and obsolescence. In formula,[18]

$$\mathcal{U}_m = \mathbb{E}\left[\sum_{H=H_0}^{\infty} e^{-rT_{H+1}^*} R\left(k_m(Q_{T_H^* + \gamma_{H+1,m}}) \right) \mathbb{1}_{\{T_{H+1,m}=T_{H+1}^*\}} - \psi \int_0^{\infty} e^{-rt} \mathbb{1}_{\{k_m(Q_t) \neq \text{OFF}\}} \, \mathrm{d}t \right] - F \tag{1.11}$$

Strategies: I focus on Markov (behavioural) strategies. Specifically, each miner m chooses the probabilities $\sigma_{m,k}(Q)$ for $k \leq Q$ of forming a blocks of size k when the mem-pool has size is Q.[19] Since $k_m(0) \in \{0\} \cup \{\text{OFF}\}$ and $k_m(Q) \in \{0, 1\} \cup \{\text{OFF}\}$, a strategy profile prescribing miners to stay always active ($\sigma_{m,\text{OFF}}(Q) = 0, \forall Q$) boils down to specifying the probabilities

$$\sigma_m = \mathbb{P}\left(k_m(Q) = 1 \mid Q \geq 1 \right)$$

Equilibrium: The equilibrium concept I employ is symmetric Markov Perfect Equilibrium (MPE). In equilibrium, $\sigma_{m,k}(Q) = \sigma_k(Q)$, $\forall m$. Of fundamental importance for the rest of the paper is the concept of *permanent mining MPE*, according to which miners are always active and fill their blocks up to their size limit, constrained by the mem-pool size if lower. Letting a Markov perfect equilibrium can be described as follows:

Definition 1.2 (Permanent mining equilibrium)**.** *A (symmetric) permanent mining equilibrium is a MPE of the mining game is such that miners are always active and fill blocks up to their limit, whenever possible. The strategy profile of a permanent mining equilibrium is such that, for all $m \in \mathcal{M}$, $\sigma_{m,0}(0) = 1$ and $\sigma_{m,1}(Q) \equiv \sigma = 1$ for all $Q \geq 1$.*

Mining tournaments

The set of strategy profiles that are part of a symmetric MPE can be determined by discarding dominated strategies and applying the one-shot deviation principle to detect profitable deviations.

Consider the actions of an active miner m after a new reference predecessor is found. If by that time the miner has not found a PoW for his block, he simply updates his ledger and

---

[18]It is possible to make the two cost components explicit by re-parametrizing the model. Assume mining generates electricity flow cost ψ' and drastic innovations occur at rate ι making the node obsolete. Then the flow cost of mining $\psi = \psi' + \iota F$.

[19]I do not consider strategies such as grim-trigger strategies or those prescribing cyclical behaviour.

continues working on extending the new longest chain. If instead the miner finds a PoW before that time, the miner transmits his block in the hope of establishing the new target block. The miner can claim a block reward only if successful in doing so. The miner's estimate of the probability of successful mining, $P(\sigma_m, \boldsymbol{\sigma_{-m}})$, depends on his block size strategy σ_m and his belief on other miners' block strategies, $\boldsymbol{\sigma_{-m}} = \{\sigma_{-m}\}_{-m \in \mathcal{M}/\{m\}}$ and can be evaluated using formula (1.7). Sometimes, this procedure does not result in closed-form expressions - e.g. when evaluating off-path beliefs resulting from a particular type of equilibrium deviation. Nevertheless, it does for the parameter configurations considered hereafter. For example, evaluated along the path of a symmetric MPE, Eq. (1.B.2) (in appendix) provides a neat result.

Lemma 1.1 (Symmetric equilibrium path)**.** *Along any symmetric equilibrium path miners have equal chances of forming the next reference predecessor block.*

$$P\left(\sigma_m = \sigma, \boldsymbol{\sigma_{-m}} = \boldsymbol{\sigma}\right) = 1/M \quad \forall \sigma \in [0,1] \tag{1.12}$$

Also, for $M = 2$, Eq. (1.B.2) determines the estimates of the successful mining probabilities for each possible equilibrium deviation in closed form.

Lemma 1.2 (Successful mining)**.** *For $M = 2$, the probability that a miner m forms the next valid block given his belief σ_{-m} on the strategy of the other miner is*

$$P(\sigma_m, \sigma_{-m}) = \frac{2\theta + (1 - \sigma_m + \sigma_{-m})\mu}{4\theta + 2\mu} \tag{1.13}$$

In particular,

$$P(1;\sigma) = \frac{2\theta + \sigma\mu}{4\theta + 2\mu} \qquad\qquad P(\sigma;\sigma) = 1/2 \quad P(0;\sigma) = \frac{2\theta + (\sigma + 1)\mu}{4\theta + 2\mu} \tag{1.14}$$

$$P(1;0) = \frac{\theta}{2\theta + \mu} \qquad P(1,1) = P(0,0) = 1/2 \quad P(0;1) = 1 - \frac{\theta}{2\theta + \mu} \tag{1.15}$$

Equilibrium block size

Here, I study under which condition a profile of equilibrium block size strategies $\boldsymbol{\sigma} = (\sigma, \sigma)$ is robust to one-shot deviations assuming miners are always active. Clearly, it makes sense to check deviations from equilibrium block sizes only for $Q \geq 1$. In this case, the expected block reward for a miner that fills a block with probability σ_m satisfies

$$\begin{aligned}
\mathbb{E}_{\sigma_m, \sigma_{-m}}\left[R\left(k_m(Q)\right)\right] &= (1 - \sigma_m)R(0)P(0;\sigma_{-m}) + \sigma_m R(1)P(1;\sigma_{-m}) \\
&= \frac{1}{4\theta + 2\mu}\left[(1 - \sigma_m)R(0)\left(2\theta + \mu(\sigma_{-m} + 1)\right) + \sigma_m R(1)(2\theta + \sigma_{-m}\mu)\right]
\end{aligned}$$

Stating explicitly that $R(1) = z(B\pi + \tau)$ and $R(0) = zB\pi$, a miner's best-response correspondence takes the following form

$$\sigma_m^* = \begin{cases} 0 & \text{for} \quad \sigma_{-m} < \frac{\pi B}{\tau} - 2\frac{\theta}{\mu} \\ [0,1] & \text{for} \quad \sigma_{-m} = \frac{\pi B}{\tau} - 2\frac{\theta}{\mu} \\ 1 & \text{for} \quad \sigma_{-m} > \frac{\pi B}{\tau} - 2\frac{\theta}{\mu} \end{cases} \tag{1.16}$$

As we can see, the best response displays strategic complementarity because is increasing in the miner's belief on his competitor's action. The unique value $\sigma = \sigma_m = \sigma_{-m}$ such that a fully mixed strategy MPE exists lies at the intersection of miners' best-responses. In other words, it is the value that makes miners indifferent between choosing any of their actions, i.e. $\mathbb{E}_{\sigma_{-m}}[R(1)] = \mathbb{E}_{\sigma_{-m}}[R(0)]$. However, such mixed-strategy equilibrium is unstable, as a small change in a miner's beliefs makes him shift away from the mixed strategy. Given the instability of the unique equilibrium with fully mixed strategies, I focus on pure strategy equilibria only.

To determine the pure strategy symmetric MPE's, notice that miners have a profitable deviation from an equilibrium prescribing $\sigma = 1$ if $\sigma_m^* = 0$ given that $\sigma_{-m} = 1$. This requires

$$\frac{\tau}{\pi} < B\left(\frac{\mu}{2\theta + \mu}\right) \tag{1.17}$$

Conversely, empty blocks equilibria are ruled out if $\sigma_m^* = 1$ given $\sigma_{-m} = 0$; that is, for

$$\frac{\tau}{\pi} > B\frac{\mu}{2\theta} \tag{1.18}$$

Conditions (1.17) and (1.18) provide a characterization of the pure strategy equilibria based on the fees-to-seigniorage ratio. If the ratio takes substantially low or high values, equilibria prescribing miners to produce empty blocks, in one case, or filled blocks, in the other, are ruled-out. For moderate values of the ratio, the game has multiple pure-strategy MPE's. Notice that higher users-to-block-size ratio ($B/1$) and block creation rate μ make condition (1.17) less stringent and have the opposite effect on Eq. (1.18); the first raises miners' seigniorage; the latter makes mining tournaments more competitive. On the contrary, a fast transmission rate θ tightens (1.17) ad softens (1.18).

Proposition 1.1 (MPE's)**.** *The equilibrium size of miners' blocks is determined by the fees-to-seigniorage ratio. Assuming $M = 2$,*

(i) For $\tau/\pi > B\mu/2\theta$, the game has a unique pure strategy MPE with $\sigma = 1$;

(ii) For $\tau/\pi < B\mu/(2\theta + \mu)$, the game has a unique pure strategy MPE with $\sigma = 0$;

(iii) For $B\mu/(2\theta + \mu) \leq \tau/\pi \leq B\mu/2\theta$, the game has two pure strategy MPE's with $\sigma = 1$ and $\sigma = 0$ and an unstable mixed strategy MPE $\sigma = B\frac{\pi}{\tau} - \frac{2\theta}{\mu}$

Notice that miners have incentive to fill blocks ($\sigma > 0$) only if $\tau > 0$. In this model, positive transaction fees are necessary for a monetary equilibrium to exist. If all blocks

are empty, validation times are infinite, so that no trader can benefit from using tokens and the cryptocurrency economy unravels. The situation is ruled-out when the fee-to-seigniorage ratio is high enough. Indeed, condition (1.18) makes mining full-blocks a *dominant strategy*, i.e. each miner prefers to mine full blocks regardless the other miners' strategies. In Section 1.6, I will show how the protocol can set the block reward in such a way that condition (1.18) holds.

In the analysis developed up to here considers an equilibrium with permanent mining; that is, $k_m(Q_t) \neq \text{OFF} \ \forall m, t$. It is important to recognize that miners have the incentive to stay active under certain conditions. In particular, a necessary condition for an equilibrium with permanent mining activity to exist is that miners make positive expected profits from each mining tournament. Since, along a symmetric MPE path, a miner can achieve the highest level of expected profits by recording a filled block, a necessary condition for mining to be viable is that the upper bound on the flow revenue exceeds the flow cost.

$$\frac{\mu}{M} z(\pi B + \tau) \geq \psi \tag{1.19}$$

On the other hand, a sufficient condition for permanent mining holds if miners prefer not to wait until recording a full block if $Q = 0$. Otherwise, in the words of Carlsten et al. (2016), an empty mem-pool causes a *mining gap*, i.e. a period of mining inactivity. Mining gaps are precluded for

$$\pi > \frac{\psi}{zB} \frac{M}{\mu} \tag{1.20}$$

Clearly the previous condition is violated for $\pi = 0$. In which case, miners only waste energy by mining empty blocks. Combining this observation with the $\tau > 0$ requirement for block filling, we can conclude that an MPE with permanent mining exists only for an interior block reward design; that is, for π and τ strictly positive.

Corollary 1.1. *A permanent mining equilibrium exists only for $\pi > 0$ and $\tau > 0$.*

Finally, an additional factor that plays a role in the activation decision is hardware deterioration, which here is neutralized by assuming no depreciation of the mining node. If mining nodes were subject to a time-increasing hazard rate of breakdown, miners would be discouraged to stay idle also by the increasing obsolescence rate.

Mem-pool and validation time distributions

Mining strategies shape the mem-pool size distribution. In particular, since valid block are transmitted approximately at rate μ, they exit the mem-pool at rate $\sigma\mu$ and enter the mem-pool at rate $\lambda = \alpha B$. The resulting stationary distribution of the mem-pool is geometric, with a probability mass function (p.m.f.) g fully determined by the (endogenous)

load factor ρ.

$$g(Q) = (1 - \rho)\,\rho^Q \qquad \text{with } \rho \triangleq \frac{\alpha B}{\sigma \mu} \in [0, 1) \tag{1.21}$$

The probabilities $g(Q), Q \in \mathbb{N}_0$ can be interpreted as the fraction of time in which the mem-pool has size Q. Thus, the quantity $1 - \rho$ is the faction of time in which the mem-pool is empty. For the steady state distribution to exist, the out-flow of transaction cannot exceed the in-flow, as imposed by Assumption 1.3.

In Section 1.4.2, the terms of trade set in DM meetings are based on the distribution of transactions' settlement time, determined by miners' criterion for picking transactions out of the mem-pool. Authors - e.g. Easley et al. (2019) - mention that miners follow the ROS (random order of service) criterion when selecting homogeneous transactions; prescribing miners to form blocks by extracting transactions uniformly at random from the mem-pool.[20] Under ROS, the number of valid blocks created (and transmitted) until a pending transaction is recorded on the blockchain is memoryless; in other words, it follows a geometric distribution. The probability that a transaction has to wait for $n \in \mathbb{N}_0$ (valid) blocks until being recorded is equal to

$$d(n) = \nu\,(1 - \nu)^{n-1}, \quad \nu \triangleq \sigma\frac{1 - \rho}{\rho}\ln(1 - \rho)^{-1} \tag{1.22}$$

The success parameter ν is the probability that a transaction is recorded in the next valid block. The derivation of the distributions (1.21) and (1.22) is presented in Appendix 1.A.1.

Miner entry

After having solved the MPE's of the mining game and computed the mem-pool and validation time distributions, I determine miners' value function to analyze their entry decision. The mining-game is connected to the cryptocurrency economy described in Section 1.3, in that each time miners earn a block reward they immediately sell it to the CM's in exchange for units of the generic good.[21]

Using DP terminology, a miner's state variables are his real balances $z_{m,t}$ and the number of pending transactions Q_t in the mem-pool. His choice variables are the size of his blocks $k_m(Q_t)$ and the activation indicator for his mining node $\chi_m(Q_t)$, set based on the equilibrium Markov strategies of the mining game. The mem-pool size Q follows the stationary distribution (1.21) at all dates, so for $\sigma = 1$, the mem-pool is filled with at least one transaction for a fraction of time $\rho = \alpha B/\mu$, and is empty for the complementary

---

[20]Quoting EOB: "when a miner builds a block he selects from the mem-pool at random instead of taking the transaction in the pool that has been waiting the longest as in a standard first-in, first-out queue."

[21] In practice, Bitcoin miners are advised to wait for the consensus chain to extend by at least 100 additional blocks (16 hours approximately) before selling a coinbase reward, so to be confident that their blocks do not become stale.

time fraction $1 - \rho = 1 - \alpha B / \mu$. Given the stationariness of Q_t, the expectation of a future block reward when miners strategize according to permanent mining equilibrium is given by

$$\mathbb{E}_Q\left[R(k_m(Q))\right] = \frac{(1 - \rho)R(0) + \rho R(1)}{M} = \frac{zB\left(\pi + \frac{\alpha}{\mu}\tau\right)}{M}$$

Taking into account that valid blocks are mined at (approximately) rate μ, we can readily formulate a miner's HJB (Hamilton-Jacobi-Bellman) equation.

Lemma 1.3. *Given a total participation of M miners, the value function of a generic miner m under permanent mining satisfies the following equations*

$$W_M^m(z_{t_0,m}) = z_{t_0,m} + W_{0,M}^m \tag{1.23}$$

$$W_{0,M}^m = \frac{1}{r}\left[\frac{\mu}{M} \times z\left(\pi B + \tau\left(\frac{\alpha B}{\mu}\right)\right) - \psi\right] - F \tag{1.24}$$

The interpretation of Eq. (1.24) is straightforward. A miner makes an up-front investment F to purchase his mining node. Afterwards, he incurs a flow cost mining ψ and receives a flow revenue from block rewards $\mu\mathbb{E}_Q\left[R(k_m(Q))\right] = \frac{Bz(\mu\pi + \alpha\tau)}{M}$.

Lemma 1.3 can be used to study miner participation under free-entry. In this case, miners join the blockchain until all mining rents are exhausted. Given that miners enter the economy with just as needed to purchase the mining node, their lifetime utility is simply given by W_0^m. so, under free entry,

$$M^* = \sup\left\{M \in \mathbb{N} : W_{0,M}^m \geq 0, W_{0,M+1}^m < 0, \ \forall m\right\}$$

Corollary 1.2 (miner participation)**.** *Under free-entry, a permanent mining equilibrium features a miner participation determined by the ratio of mining revenues to costs,*

$$M = \lceil\tilde{M}\rceil \quad \tilde{M} \triangleq \frac{Bz\left(\mu\pi + \alpha\tau\right)}{\psi + rF} \tag{1.25}$$

Traders' trust in the blockchain depends on its security level, which increases with miner participation. To ensure a secure participation level, traders require $M \geq \underline{M}$. The problem of setting the optimal block reward to induce a given level of miner participation is studied in Section 1.6, where I will present a simple design for $\underline{M} = 2$.

1.4.2 Trade

In this section, I analyse the partial equilibrium model of trade that determines buyers'
optimal (cryptocurrency) portfolio choice and traders' DM bargaining conditions. Sellers
are passive except when setting the terms of trade. The HJB equations that solve traders'
DP problems are again presented in Appendix 1.B.

Buyers

Each buyer can continuously trade in her CM if not busy in a DM meeting. A buyer b
participating to her CM with a stock of real balances $z_{b,t}$ chooses to supply $x = -(z^* - z_{b,t})$
units of the generic good to adjust her real balances to a desired level z^*. Also, at a Poisson
rate α, she engages in a DM meeting with a (randomly drawn) seller s. If the seller is
available, CM trade is interrupted by a DM meeting in which b trades her tokens for
units of the special good produced by s. The number of sellers is large enough to ensure
that buyers almost always meet an available seller, so the total rate of DM meetings is
$\lambda = \alpha B$.

In what follows, buyers can perfectly observe the blockchain and adopt the portfolio
strategy of maintaining their real balances at the desired level z^* until they enter a DM
meeting. To do so, buyers make a tiny adjustment in their token portfolio each time
a block is produced to compensate for the inflation shock caused by a miner selling the
associated coinbase. To be precise, if a buyer has z real balances before a coinbase reward
is sold to the CM's, their value drops to $z(1 - \pi)$ afterwards. Thus, to maintain a desired
amount z^*, she has to acquire a value of $z^*\pi$ from her CM.[22]

According to a DP formulation, the state variable of a buyer is given by her real
balances $z_{b,t}$ and her control variable by the desired amount of balances z^* held until the
next DM meeting occurs, or analogously by her generic good supply $x = -(z^* - z_{b,t})$. The
CM and DM value functions for a buyer, $W^b(z_{b,t})$ and $V^b(z_{b,t})$, given her real balances
$z_{b,t}$ obey the HJB equations formulated in the next lemma.

Lemma 1.4. *A representative buyer's value function is given by*

$$W^b(z_{b,t}) = z_{b,t} + W_0^b$$

$$rW_0^b = -z^* \left(r + \mu\pi\right) + \alpha \left[V^b(z^*) - W^b(z^*)\right] \tag{1.26}$$

Intuitively, buyers bear the cost of adjusting real balances to the optimal quantity z^*
together with the cost of holding tokens without spending them originating from inflation
and discounting. At rate α, buyers use their optimal amount of tokens to obtain a capital

---

[22]It would be more natural to assume buyers cannot observe the blockchain and let them set a
continuous adjustment rule based on expected inflation.

gain from DM trade. The DM value function $V^b(z^*)$ depends on the agreed bargaining terms between buyer and seller.

Sellers

Sellers enter the CM with real balances $z_{s,t}$ and a set of pending transactions waiting for validation on the blockchain and engage in DM according to a meeting rate α. To keep the model tractable, I assume that sellers are too busy to keep track of the number of the number of pending transactions they are waiting to receive and discount each of them independently. Hence they neglect the congestion effect caused by multiple pending transactions. As long meetings with buyers are rare enough at individual seller level, this assumption has negligible implications for the trade equilibrium. Sellers are passive except for selling their tokens, when available, and defining the terms of DM trade with buyers.

The value of a pending DM transaction is eroded by the depreciation caused by discounting and inflation until it is settled. For this reason, a transaction involving $z_{s,t}$ real tokens waiting for validation is discounted using a specific *validation discount factor* $\Omega \in [0,1]$ that accounts for the effects of inflation and time discounting during the transaction settlement period. The resulting present value of a pending DM transaction is thus $z_{s,t}\Omega$.

The validation discount factor can be determined through the following observations. First, the discount factor that applies to a pending transaction waiting to be recorded in the next valid block is the product of a term resulting from temporal discounting, $\frac{\mu}{r+\mu}$, (see Eq. (1.A.5)) and an inflation adjustment term, $1 - \pi$, resulting in the expression

$$\frac{\mu}{r+\mu}(1-\pi) \tag{1.27}$$

Since a pending transaction is discounted by (1.27) each time the transaction waits for an additional valid block to be recorded, and given that the number of valid blocks posted until a transaction is recorded on the blockchain is distributed according to the geometric p.m.f. $d(n)$ from Eq. (1.22), the validation discount factor Ω follows from the next expectation:

$$\Omega \triangleq \mathbb{E}_n\left[\left(\frac{\mu}{r+\mu}(1-\pi)\right)\right]^{n+1} = \frac{(1-\pi)\mu\nu}{r+\mu\left(\nu+\pi(1-\nu)\right)} \qquad \text{with } \Omega_\pi < 0,\ \Omega_{\pi\pi} > 0,\ \Omega_{\pi,\nu} < 0 \tag{1.28}$$

where the block inclusion probability ν is defined in Eq. (1.22). Now I am ready to present the representative seller's value function.

Lemma 1.5. *The value function of a seller at the time he records a pending transaction is given by*

$$W^s(\Omega z_{s,t}) = \Omega z_{s,t} + W_0^s \quad with \quad rW_0^s = \alpha\left(V^s - W_0^s\right) \tag{1.29}$$

Where Ω is the validation discount factor determined in (1.28)

Intuitively, the seller's expected enjoyment from generic good consumption once his pending tokens are settled and immediately cashed-out to a CM amounts to $\mathbb{E}(x) = z_{s,t}\Omega$. In the meantime, the seller continues to trade at rate α with buyers.

In practice and in Chiu and Koeppl (2019), before delivering the special good to a buyer, sellers wait for the consensus chain to grow until the block recording her payment is deep enough inside it, so that the blocks built on top of it count as "confirmations" of the block's validity. Realistic terms of trade should also cover the *depth* of the valid block including the payment, where "depth" is intended as the distance between a block and the terminal block of the consensus chain. In BTC, the praxis is to wait for the consensus chain to grow by 6 blocks before accepting a payment, so that it takes one hour for a transaction to be considered safe. In Ethereum (ETH), sellers are advised to wait for 30 confirmations but block creation is considerably faster, so that it takes on average 6 minutes for a transaction to be safely considered as valid. Sellers accepting Bitcoin Cash (BCH) payments are suggested to wait for 15 confirmations so that transactions become valid, in expectation, after 2 hours and a half.

DM bargaining

When entering the DM, a buyer-seller pair *b-s* meets to agree on the terms of DM trade $(y, p(y))$ specified by an amount of special good units produced by s and a price paid by b. Upon meeting, traders observe the current token prices and b reveals to s his real balances z. Since carrying liquidity is costly, b leaves the DM without tokens, so that $p(y) = z^*$ holds, i.e. the cash-in-advance constraint is binding. To facilitate exposition, in this section I let $z^* \equiv z$ and denote as $y(z)$ the amount of special good produced by the seller given his knowledge of the buyer's real balances.

Since digital wallets are programmed to charge a transaction fee τ, the seller receives only a part $z(1-\tau)$ of the value transferred by the buyer. The remaining $z\tau$ goes to the miner that records the transaction on the blockchain.

Right after setting the trade terms but before re-joining their respective CM's, the buyer enjoys utility $u(y(z))$ from consuming the special good and sends her tokens to the seller, while the seller spends a cost y to produce the special good and records a pending transaction of value $z(1-\tau)$. It follows that

$$V^b(z) = u\left(y(z)\right) + W_0^b \tag{1.30}$$

$$V^s = -y(z) + W^s\left(\Omega z(1-\tau)\right) \tag{1.31}$$

I determine the terms of trade using Kalai bargaining, according to which traders split

the trade surplus in proportion to their bargaining power. To do so, notice that traders' outside options are $W^b(z)$ and W_0^s, hence their trading surpluses follow the equations

$$V^b(z) - W^b(z) = \underbrace{u\left(y(z)\right) + W_0^b - W^b(z)}_{= u\left(y(z)\right) - z} \quad \text{and} \quad V^s - W_0^s = \underbrace{-y(z) + W^s\left(\Omega z(1-\tau)\right) - W_0^s}_{= -y(z) + \Omega z(1-\tau)}$$

$$(1.32)$$

Thus, letting β denote buyers' bargaining weight, Kalai bargaining implies

$$u(y) - p(y) = \frac{\beta}{1-\beta}\left(-y + p(y)(1-\tau)\Omega\right) \quad \text{equivalent to} \quad p(y) = \frac{(1-\beta)u(y) + \beta y}{(1-\beta) + \beta(1-\tau)\Omega}$$

$$(1.33)$$

Plugging the resulting special good price back into the trade surplus expressions (1.32), we can see that traders share the (adjusted) total surplus in proportion to their bargaining power, so that the DM capital gain buyer and seller

$$V^b(z) - W^b(z) = \beta \frac{u(y(z))(1-\tau)\Omega + y(z)}{(1-\beta) + \beta(1-\tau)\Omega} \quad V^s - W_0^s = (1-\beta)\frac{u(y(z))(1-\tau)\Omega + y(z)}{(1-\beta) + \beta(1-\tau)\Omega}$$

Applying the formulae for the trade surpluses and special good price to these DM value functions, traders' DP equations now take explicitly into account the role of trade frictions, so that I can solve for buyers' optimal real balance portfolio z^*.

Lemma 1.6 (special good production and optimal token portfolio). *The optimal amount of special good produced by sellers and the optimal token portfolio held by buyers (in real terms) are determined as follows:*

$$y^* : \frac{1}{\alpha}\left(r + \mu\pi\right) \underbrace{\leq}_{= \text{ if } y^* > 0} \frac{u'(y^*)}{p'(y^*)} - 1, \quad z^* = p(y^*) \tag{1.34}$$

The optimal token portfolio is determined by equating costs and benefits of carrying tokens. The cost of increasing marginally real balances is the sum of the marginal discounting r and inflation cost $\mu\pi$ paid to maintain the optimal amount of real balances while keeping-up with inflation. Both costs are incurred for a period of time of expected length $1/\alpha$. The benefit of holding marginally more liquidity (tokens in this case) is expressed by the right-side of inequality (1.34) and measures the sensitivity of the DM capital gain with respect to real balances. The LW literature defines this marginal effect as the "liquidity premium" (see Choi and Rocheteau (2020b) or Lagos et al. (2014) for some examples of usage).

As previously anticipated, for a positive inflation rate buyers do not carry more tokens than those needed in the DM. The opposite happens for a (sufficiently) deflationary cryptocurrency whose tokens gain value over time. This feature makes a deflationary cryptocurrency more suited to be used as a safe-heaven asset rather than as a means of payment.

An explicit formula for y^* is easily obtained from Eqs. (1.33) and (1.34) under the assumption that buyers have full bargaining power ($\beta = 1$). Making explicit the utility function (1.1), we have

$$y^* = y(z^*) = \frac{\alpha(1-\tau)\Omega}{r + \alpha + \mu\pi} - \frac{1}{\eta} \tag{1.35}$$

$$z^* = \frac{\alpha}{r + \alpha + \mu\pi} - \frac{1}{\eta(1-\tau)\Omega} \tag{1.36}$$

Buyers' desired level of special good consumption y^* is decreasing in the transaction fee rate and congestion (lower Ω) owing to worse contractual terms. Notice that a large enough value of η ensures that $y^* > 0$.

1.5 Monetary equilibrium

In this section, I combine the partial equilibrium models of Section 1.4 and determine the monetary (general) equilibrium of the economy.

Equation (1.33) can be used as an asset-pricing identity by decomposing real balances into a number of tokens held and tokens' price: $p(y^*) = z^* = \phi_{i,t}a_{b,t}$, where i is the index of the CM in which b operates. CM prices follow from aggregating the condition (1.33) over the CM's buyer population.

$$\phi_{i,t} = \frac{p(y^*)}{A_{i,t}}B_i = \quad \text{for } i = 1, 2$$

The Market clearing condition (1.2) together with the stationary property (1.3) and the inflation schedule (1.9) pin-down the dynamic evolution of a DM price for a given initial amount of tokens in circulation A_{i,t_0} $\forall i$ as follows. Letting b_t indicate the number of valid blocks produced by time t,

$$\mathbb{E}\left(\phi_{i,t}\right) = \frac{B}{2}\frac{p(y^*)}{A_{i,t_0}}\frac{\mathbb{E}\left(\phi_{i,t} \mid \phi_{i,t_0}\right)}{\phi_{i,t_0}} \quad \text{with} \quad \frac{\mathbb{E}\left(\phi_{i,t} \mid \phi_{i,t_0}\right)}{\phi_{i,t_0}} = \mathbb{E}\left[(1-\pi)^{b_t-b_{t_0}}\right] = e^{-\pi\mu(t-t_0)}$$

$$= \frac{B}{2}\frac{p(y^*)}{A_{i,t_0}}e^{-\pi\mu(t-t_0)} \tag{1.37}$$

A complete derivation of the previous formula is contained in Appendix 1.B. Expected prices fall exponentially at the rate of seigniorage and depend positively on tokens' utilization through the special good price.

The general monetary equilibrium of the economy puts together all elements determined so far and closes the model.

Definition 1.3 (Cryptocurrency Equilibrium). *Given the design parameters (μ, π, τ) and initial tokens in circulation A_{i,t_0} for $i = 1, 2$, a cryptocurrency (monetary) equilibrium*

with permanent mining is a collection of value functions (W_0^m, W_0^b, W_0^s), (1.23,1.26,1.29) optimal controls (σ^*, z^*, y^*) (1.16, 1.34, 1.35), prices $\left(p(y), \left(\phi_{i,t}\right)_{i=1,2}\right)$ (1.33,1.37) and distributions $(g(\cdot), d(\cdot))$ (1.21,1.22) such that (i) all agents make optimal decisions (ii) market clearing (1.2) holds in each CM (iii) the validation discount factor applied to DM transactions is determined by equation (1.28). (iv) miners play a permanent mining equilibrium and entry the economy according to expression (1.25).

This equilibrium is a standard forward-looking monetary equilibrium with policy parameters linked to the blockchain's internal functioning. Monetary policy is determined by the amount of tokens created by coinbase transactions and the block creation rate. Record-keeping is performed by miners, provided that they have the incentive to do so.

The existence of a cryptocurrency equilibrium requires a contained level of congestion. If congestion gets out of control, settlement times become huge, so that sellers charge a prohibitive price for the special good and no buyer can benefit from trading in the DM market, the only reason for them to hold tokens in this model. In the extreme case in which miners refuse to fill blocks, congestion becomes infinite. It is therefore essential that the protocol is designed in such a way to incentivize mining.

1.6   Optimal cryptocurrency design

In this section, I take a normative standpoint and study the optimal design of a cryptocurrency. Specifically, I study the problem of a blockchain protocol designer who sets inflation, block rate, and transaction fees to elicit a secure level of miner participation and incentivize miners to fill blocks. The designer, or social planner (SP), can be thought of as a software developer designing a modification of the Bitcoin source code.

In a standard optimal taxation problem, the government can dictate the supply of public goods. In blockchain design instead, miners' behavior is guided by their incentives, so the protocol sets the policy instruments (μ, π, τ) to induce miners in providing a level of public goods $(M, k(Q))$. Miner participation M raises blockchain security, a pure public good, but also creates a welfare loss caused by the PoW energy costs.[23] On the other hand, processing capacity (or mining throughput) resulting from miners' block size strategies $k(Q)$ is a common-pool good subject to congestion.

The planner's problem can be formalized as the maximization of agents' aggregate utility subject to mining and trade incentive compatibility constraints. Maintaining the assumption that buyers have full bargaining power simplifies the welfare expression since $W_0^{s,0} = 0$. Furthermore, the free-entry condition is such that miner participation is

---

[23]Current estimates from https://digiconomist.net/bitcoin-energy-consumption indicate that Bitcoin annual energy consumption accounts for approximately 0.21% of the world's total, producing a carbon footprint comparable to that of a small country - around 33 CO2 megatons. Several protocol proposals aim to address the sustainability issue associated with Bitcoin PoW.

the highest admissible value that provides miners a non-negative surplus. If M where continuous, free-entry would set mining rents to zero. However, under the assumption that M is discrete, miners can still earn a positive rent due to a mere integer rounding. So, for simplicity, I will ignore the rounding issue and treat miners as if they earn no surplus, i.e. setting $W_0^{m,0} = 0$. As a result of these simplifications, the planner simply maximizes per-buyer surplus.

$$W^{SP} = \max_{\mu,\pi,\tau} W_0^b, \quad W_0^b = \frac{\alpha}{r}\left(u\left[y(z^*)\right]\right) - z^*\left(1 + \frac{\mu\pi}{r}\right) \tag{1.38}$$

$$\text{subject to } M = \underline{M} \text{ and optimal } (\sigma^*, y^*, z^*)$$

Notice that eliciting miner participation beyond $\underline{M}$ creates only a dead-weight loss, so letting miners only up to the secure level of participation is optimal in any design. Also, since the mixed-strategy mining MPE is not a reasonable prediction of mining behaviour, the planner has to preclude miner inactivity at any point in time to make sure that miners perform their record-keeping task when there are pending transactions and continue building the blockchain keep it secure when the mem-pool is empty. In other words, the planner implements a permanent mining equilibrium.

1.6.1 Contribution to the fees-versus-seigniorage debate

In Chiu and Koeppl (2019), an optimally designed cryptocurrency rewards miners only with seigniorage. CK's result follows from the observation that inflation is paid on a larger "tax base" than transaction fees, so the designer can guarantee miners a given amount revenues by charging commodity traders at a lower rate than one implied by a design that also uses fees. However this argument ignores the effects of inflation on the rate of per-block transactions that miners record and hence on the speed at which payments are processed.

In particular, with endogenous block size, raising inflation reduces the size of blocks formed by miners, which in turn produces the welfare-reducing effect of making settlement times of pending transactions longer. In the extreme case in which miners produce only empty blocks, settlement times are infinite, so the cryptocurrency infrastructure cannot sustain a monetary equilibrium.

In this model, a pure seigniorage reward causes this scenario to occur. Nevertheless, the protocol can avoid such situation by setting a lower bound on fees so to ensure that miners take charge of their role of record-keepers. The safest scenario that the protocol can enforce is the one in which miners have a dominant strategy to play a permanent (maximal) mining equilibrium. In other words, each miner finds optimal to fill blocks regardless what other miners do. If this is the case, a monetary equilibrium is implementable in dominant mining strategies.

Definition 1.4. *A policy (μ, π, τ) implements a mining equilibrium $\boldsymbol{\sigma}(\mu, \pi, \tau)$ in dominant strategies if and only if*

$$\mathbb{E}_{\sigma(\mu,\pi,\tau),\boldsymbol{\sigma}_{-m}} \left[R\left(k_m(Q)\right) \mid Q \right] \geq \mathbb{E}_{\sigma'_m,\boldsymbol{\sigma}'_{-m}} \left[R\left(k_m(Q)\right) \mid Q \right], \quad \forall \sigma'_m, \boldsymbol{\sigma}_{-m}, Q$$

Dominant strategy implementation of maximal mining requires the fee-to-seigniorage ratio to satisfy Proposition 1.1-(i). This combined with the activity condition (1.20) and security condition $M \equiv \lceil M^* \rceil = \underline{M}$, with equilibrium miner entry given by (1.25). These constitute the constraints that the planner has to satisfy when choosing policy parameters optimally. To summarize, the planner maximizes (1.38) subject to

$$\frac{\tau}{\pi} \geq \qquad\qquad \frac{\mu B}{2\theta} \qquad \text{Record-keeping constraint} \qquad \text{(KEP)}$$

$$2 \geq \frac{Bz^*(\pi,\tau)(\mu\pi + \alpha\tau)}{\psi + rF} \qquad > 1 \qquad \text{Security constraint} \qquad \text{(SEC)}$$

$$\pi \geq \qquad\qquad \frac{2\psi}{\mu B} \qquad \text{Activity constraint} \qquad \text{(ACT)}$$

The general design problem is complicated and beyond the scope of the current version of this paper that only analyzes record-keeping incentives with a two-miner population. Yet, we can already provide a simple optimal design with the analysis developed so far. In particular, notice that if a secure participation $\underline{M} = 2$ can be achieved with the lowest admissible value (π^*, μ^*) implied by the above constraints, that is an optimal design. Even though this is a simplistic approach to the general problem, it is already instructive as it shows an example where the optimal level of fees is positive. The following proposition presents a simple closed-form expression for such minimal feasible rates of fees and seigniorage as well the condition under which miner secure entry is achieved and hence the design is optimal.

Proposition 1.2. *For $\underline{M} = 2$, the lowest rates (π, τ) that can implement a monetary equilibrium in dominant mining strategy $(\sigma = 1)$ are*

$$\pi = \frac{2\psi}{\mu B} \qquad \tau = \frac{\psi}{\theta} \tag{1.39}$$

For fixed μ, these rates constitute an optimal design if the security constraint is (SEC) is satisfied at those values.

Notice that both seigniorage and transaction fee rates are increasing in the energy cost of mining, but the fee rate decreases in the block propagation rate, which raises miner chances to record blocks and cash-in the fees associated to them, while the inflation rate decreases in the size of the buyer population at which cryptocurrency deposits will be charged.

1.6.2 Discussion

Proposition 1.2 assumes that the security requirement $\underline{M}$ is independent of the policy instruments mix. In CK instead, higher fees increase the secure (honest) miners' participation requirement because, by reversing transactions with DS attacks, fraudulent miners - that in CK are also buyers - reverse to themselves the value of transaction fees as well. However, in practice, transaction fees are small relative to the capital gain that makes a fraud worthwhile and have mild effects on miners' incentives to misbehave. For this reason, it is realistic to consider as negligible the impact of τ on $\underline{M}$.

Finally, an alternative but less robust argument for justifying reducing seigniorage in favour of transaction fees is based on the "hot potato effect" of inflation. According to this argument, higher seigniorage (and inflation), raises the velocity of money, i.e. induces traders to spend their tokens at a faster rate, creating a negative congestion externality. Analyzing this argument requires modelling buyers' endogenous intensity of search for sellers, possibly adapting Ennis (2009) model to continuous time cryptocurrency economy studied so far.

1.7 Conclusion

I presented a model that describes the plurality of interactions involved in a cryptocurrency economy. Miners play the essential role of record-keepers by recording blocks of transactions on their copies of the cryptocurrency blockchain. They also contribute to the blockchain security and resolve inconsistencies among their ledger copies by following the Longest-Chain-Rule, indicating them a unique branch of the ledger to follow.

Sellers receive transactions from buyers only after they are recorded by miners on the blockchain. As a result, transaction settlement is stochastic and depends on miners' block size strategy. Since larger blocks imply a short transaction settlement period, traders are always better-off if miners choose to create large blocks; however, miners do not always benefit from doing so. Indeed, miners increase the size of their blocks only if the higher transaction fees compensate for the higher risk that their blocks become stale due to their slow transmission time caused by having a large dimension.

Small blocks lead to a higher load of pending transactions on the blockchain and worsen the terms of cryptocurrency trade. In the extreme case in which miners only create empty blocks, transactions never get processed, and a monetary equilibrium in which cryptocurrency trade takes place is not plausible. A clever blockchain design can avoid this scenario by setting high enough transaction fees relative to seigniorage, thereby providing miners with the incentive to fill blocks.

The analysis developed in this paper studied in detail the basic mining trade-off involving block reward and the risk of blocks becoming stale caused by slow block transmission.

Further research could extend the model developed in this paper by introducing verification times and SPV mining. In practice, each time a miner receives a block he has to spend a *verification time* to check the block's validity. Afterwards, if the verification has a positive outcome, he updates his ledger and mem-pool with the information contained in the verified block, whereas in the contrary case - e.g. if the block contains double-spent transactions or an invalid PoW - he forks-it out in a stale branch. A miner cannot safely mine new non-empty blocks while preforming a verification, as they can cause a double-spending with the transactions included in the block under verification, but faces no risk in mining empty blocks while verification is still incomplete. The practice of creating empty blocks skipping full verification is known as Simplified Payment Verification (SPV) mining. The practice has the positive effect of precluding mining-gaps and thus rises mining revenues and in turn the level of blockchain security against malicious attacks. These mining incentives are particularly important when miners' mem-pools are empty ($Q = 0$). Nevertheless, SPV mining is a form of free-riding when miners' mem-pools are non-empty ($Q > 0$) and allows record-keeping errors to remain unnoticed for long, amplifying their detrimental effects when finally revealed.

Finally, by letting cryptocurrency prices vary across CM's, the model can be used to address traders' incentives to engage in speculation.

Appendices of Chapter 1

1.A Additional derivations

1.A.1 Distributions of mem-pools' size and validation times

Mem-pools

The steady state distribution of the mem-pool balances expected in-flows and out-flows of transactions. For $Q_t > 0$, miners receive transactions at rate $\lambda \equiv \alpha B$ and process transactions at rate $\sigma\mu$. If $Q_t = 0$ instead, miners have no transaction to process. Let $\dot{Q}_t \triangleq \partial Q_t / \partial t$. We have

$$\dot{g}(0) = \mu\sigma g(1) - \lambda g(0) \tag{1.A.1}$$

$$\dot{g}(Q) = \lambda g(Q-1) + \mu\sigma g(Q+1) - (\lambda + \mu\sigma)\, g(Q) \qquad \text{for } Q > 0 \tag{1.A.2}$$

$$\text{with } \sum_{Q=0}^{\infty} g(Q) = 1 \tag{1.A.3}$$

Setting $\dot{g}(Q) = 0 \; \forall Q$ yields $g(Q) = \rho g(Q-1)$, where $\rho = \lambda/\sigma\mu$ denotes the load factor, implying $g(Q) = \rho^Q g(0)$ for all $Q \geq 0$. Applying the normalization (1.A.3) we obtain the stationary probability of the mem-pool being empty,

$$\sum_{Q=0}^{\infty} \rho^Q g(0) = \frac{g(0)}{1-\rho}, \text{ so, } g(0) = 1 - \rho$$

Now we can see that the mem-pool distribution is geometric with p.m.f. $g(Q) = (1-\rho)\rho^Q$
∎

Remark: In a permanent mining equilibrium ($\sigma = 1$) the mining game defined in Section 1.4.1 switches across mem-pool states according to the following transition rate matrix

$$
\begin{array}{c c}
 & \begin{array}{cccc} (Q,H) & (Q+1,H) & (Q,H+1) & (Q-1,H+1) \end{array} \\
\begin{array}{c} (Q,H), \; Q=0 \\ (Q,H), \; Q>0 \end{array} &
\left[\begin{array}{cccc}
1-\lambda-\mu & \lambda & \mu & 0 \\
1-\lambda-\mu & \lambda & 0 & \mu
\end{array} \right]
\end{array}
$$

Validation time

Suppose miners form blocks picking transactions uniformly at random from their mem-pools and each seller has at most one pending transaction. The probability that a pending transaction directed to seller s is recorded in a block is $1/(1 + Q_{-s})$, where Q_{-s} is the number of pending transactions not yet recorded and directed towards other sellers. Since seller s takes into account that Q_{-s} is geometrically distributed with p.m.f. $g(Q)$ (from Eq. (1.21)), she computes the probability that her transaction will be included in the next block using the following formula:

$$\sum_{Q_{-s}=0}^{\infty} g(Q_{-s}) \frac{1}{1+Q_{-s}} = (1-\rho) \sum_{Q_{-s}=0}^{\infty} \frac{\rho^{Q_{-s}}}{1+Q_{-s}} = \frac{1-\rho}{\rho} \left(\sum_{Q=1}^{\infty} \frac{\rho^Q}{Q} \right) = \frac{1-\rho}{\rho} \ln(1-\rho)^{-1}$$
$$(1.A.4)$$

where the last equality follows from the Maclaurin expansion of $\ln(1-\rho)^{-1}$. Since miners fill blocks with probability σ, it follows that a transaction's validation time in block units is geometrically distributed with p.m.f. $d(n) = (1-\nu)^{n-1}\nu$, in which the block inclusion parameter ν satisfies $\nu = \sigma\frac{1-\rho}{\rho}\ln(1-\rho)^{-1}$ ∎

1.A.2 Properties of agents' value functions

To follow the main proofs in Appendix 1.B and part of the analysis contained in the main body of the paper, it is worth to keep in mind some properties of value functions driven by an underlying Poisson process and Poisson processes themselves.

Exponential random variables:

Let $\{T_j\}_{j\in\{1,2,\dots,J\}}$ denote a collection of exponential random variables with with rates $\{\alpha_j\}_{j\in\{1,2,\dots,J\}}$. Their minimum $T = \min\{T_j\}_{j\in\{1,2,\dots,J\}}$ is again exponentially distributed with rate $\alpha = \sum_{j\in\{1,2,\dots,J\}}\alpha_j$. Moreover, since the Laplace-Stieltjes transform of the exponential random variable T_j is given by $\mathbb{E}\left(e^{-rT_j}\right) = \frac{\alpha_j}{r+\alpha_j}$, so we also have that

$$\mathbb{E}\left(e^{-rT}\right) = \alpha/(r+\alpha) \tag{1.A.5}$$

Stochastic upper integration limit:

Let T_j denote an exponential RV with rate α_j. The following identity holds:

$$\mathbb{E}_{T_j}\left(\int_0^{T_j} e^{-rt} f(t)\,\mathrm{d}t\right) = \int_0^\infty e^{-(r+\alpha_j)t} f(t)\,\mathrm{d}t \tag{1.A.6}$$

To see this, integrate the left-hand side of Eq. (1.A.6) by parts, setting $v(T_j) = -e^{-\alpha_j T_j}$ and $u'(T_j) = e^{-rT_j} h_{T_j}$ to obtain

$$\mathbb{E}_{T_j}\left(\int_0^{T_j} e^{-rt} f(t)\,\mathrm{d}t\right) = \int_0^\infty \alpha_j e^{-\alpha_j T_j}\left(\int_0^{T_j} e^{-rt} f(t)\,\mathrm{d}t\right)\mathrm{d}T_j$$

$$= \int_0^\infty v'(T_j)u(T_j)\,\mathrm{d}T_j = \underbrace{\Big[v(T_j)u(T_j)\Big]_0^\infty}_{=0} - \int_0^\infty v(T_j)u'(T_j)\,\mathrm{d}T_j$$

$$= \int_0^\infty e^{-(r+\alpha_j)T_j} f(T_j)\,\mathrm{d}T_j \equiv \int_0^\infty e^{-(r+\alpha_j)t} f(t)\,\mathrm{d}t \quad \blacksquare$$

Switching Poisson states:

Suppose a value function W^0 can switch from state 0 to states $j \in \{1,2,\dots,J\}$ at a Poisson rates $\alpha_{0j} \equiv \alpha_j$. Then, as we saw earlier in this section, the switching-time T from state 0 to a next generic state is exponentially distributed with cumulative transition rate α and the probability that the value function will enter a given state j next is α_j/α. Therefore,

$$W^0 = \mathbb{E}_{T,j}\left(e^{-rT}W^j\right) = \frac{\alpha}{r+\alpha}\left(\sum_{j\in\{1,2,\dots,J\}} W^j\frac{\alpha_j}{\alpha}\right) = \frac{1}{r+\alpha}\sum_{j\in\{1,2,\dots,J\}} \alpha_j W^j$$

It follows that

$$rW^0 = \sum_{j\in\{1,2,\dots,J\}} \alpha_j\left(W^j - W^0\right) \quad \blacksquare \tag{1.A.7}$$

1.B Omitted proofs

Proof of Lemma 1.1. If $\gamma_{H+1,m} + \epsilon_{H+1,m} < \min_{m'} \left\{ \gamma_{H+1,m'} + \epsilon_{H+1,m'} \right\}$, miner m's transmission time is fast enough to let him establish the next reference predecessor block and collects the block reward. If not, m's block gets forked-out by all other miners. For now, let $\gamma_{H+1,m} \equiv \gamma_m$ and $\epsilon_{H+1,m} \equiv \epsilon_m$ for a given $m \in \mathcal{M}$. The probability $\tilde{P}_{M-1}(T; \boldsymbol{\sigma_{-m}} = \boldsymbol{\sigma})$ that a block with update time $T_{H+1,m} \equiv T$ is faster than the contemporaneous $M-1$ blocks, evaluated using the belief $\boldsymbol{\sigma_{-m}} = \boldsymbol{\sigma}$ on other miners' block strategies, satisfies the following relationships

$$\tilde{P}_{M-1}(T; \boldsymbol{\sigma_{-m}} = \boldsymbol{\sigma}) = \mathbb{P}\left(\min_{m'} \left\{ \gamma_{m'} + \epsilon_{m'} \right\} > T \right) = \left(1 - F_{\gamma_{m'}+\epsilon_{m'}}(T) \right)^{M-1}$$

$$= \left[e^{-T(\mu/M)} + \sigma \frac{\mu/M}{\theta - (\mu/M)} \left(e^{-T(\mu/M)} - e^{-T\theta} \right) \right]^{M-1} \qquad (1.B.1)$$

The unconditional belief of successful mining $P(\sigma_m, \boldsymbol{\sigma_{-m}})$ is obtained from expression (1.B.1) by integrating-out the marginal density of T, so that

$$P(\sigma_m; \boldsymbol{\sigma_{-m}}) = \int_0^\infty \tilde{P}(t, \boldsymbol{\sigma_{-m}}) f_{\gamma_m+\epsilon_m}(t) \, \mathrm{d}t \qquad (1.B.2)$$

Now, integrating Eq. (1.B.2) by parts,

$$\int_0^\infty P_1(t)^{M-1} f(t) \, \mathrm{d}t = \left[P_1(t)^{M-1} \left(1 - P_1(t) \right) \right]_{t=0}^{t \to \infty} + (M-1) \int_0^\infty P_1(t)^{M-2} \left(1 - P_1(t) \right) f(t) \, \mathrm{d}t$$

$$= \underbrace{\left[P_1(t)^{M-1} - P_1(t)^M \right]_{t=0}^{t \to \infty}}_{=0} + (M-1) \int_0^\infty P_1(t)^{M-2} f(t) - P_1(t)^{M-1} f(t) \, \mathrm{d}t$$

$$\int_0^\infty P_1(t)^{M-1} f(t) \, \mathrm{d}t = (M-1) \left[\int_0^\infty P_1(t)^{M-2} f(t) \, \mathrm{d}t \right] - (M-1) \left[\int_0^\infty P_1(t)^{M-1} f(t) \, \mathrm{d}t \right]$$

Collecting the identical integrals,

$$\int_0^\infty P_1(t)^{M-1} f(t) \, \mathrm{d}t = \frac{M-1}{M} \left[\int_0^\infty P_1(t)^{M-2} f(t) \, \mathrm{d}t \right] \qquad (1.B.3)$$

Expression (1.12) follows immediately from Eq. (1.B.3) proceeding by induction on M with base step $M = 2$ ∎

Proof of Lemma 1.2. From Eq. (1.B.2) the probability we are looking for is the result

of the following integral

$$P(\sigma_m, \sigma_{-m}) = \int_0^\infty \left[e^{-T\frac{\mu}{2}} + \sigma_{-m}\frac{\mu}{2\theta - \mu}\left(e^{-T\frac{\mu}{2}} - e^{-T\theta}\right) \right]$$

$$\times \left[\sigma_m \frac{\theta\mu}{2\theta - \mu}\left(e^{-T\frac{\mu}{2}} - e^{-T\theta}\right) + \frac{\mu}{2}(1 - \sigma_m)e^{-T\frac{\mu}{2}} \right] dT \quad (1.B.4)$$

Expand the product in the previous expression and integrate each part separately. We have that

$$\left(\sigma_m \frac{\theta\mu}{2\theta - \mu}\right) \int_0^\infty e^{-T\frac{\mu}{2}}\left(e^{-T\frac{\mu}{2}} - e^{-T\theta}\right) dT$$

$$= \left(\sigma_m \frac{\theta\mu}{2\theta - \mu}\right) \frac{2\theta - \mu}{\mu(2\theta + \mu)}$$

$$= \frac{\sigma_m\theta}{2\theta + \mu}$$

$$\frac{\mu}{2}(1 - \sigma_m) \int_0^\infty e^{-T\mu} dT = \frac{1 - \sigma_m}{2}$$

$$\left(\sigma_{-m}(1 - \sigma_m)\frac{\mu^2}{2(2\theta - \mu)}\right) \int_0^\infty e^{-T\frac{\mu}{2}}\left(e^{-T\frac{\mu}{2}} - e^{-T\theta}\right)$$

$$= \left(\sigma_{-m}(1 - \sigma_m)\frac{\mu^2}{2(2\theta - \mu)}\right) \frac{2\theta - \mu}{\mu(2\theta + \mu)}$$

$$= \frac{(1 - \sigma_m)\sigma_{-m}\mu}{2(2\theta + \mu)}$$

$$\left(\sigma_{-m}\sigma_m \frac{\mu^2\theta}{(2\theta - \mu)^2}\right) \int_0^\infty \left(e^{-T\frac{\mu}{2}} - e^{-T\theta}\right)^2 dT$$

$$= \left(\sigma_{-m}\sigma_m \frac{\mu^2\theta}{(2\theta - \mu)^2}\right) \frac{(\theta - \mu/2)^2}{\theta\mu(\theta + \mu/2)}$$

$$= \frac{\sigma_{-m}\sigma_m\mu}{2(2\theta + \mu)}$$

Summing up the four factors we obtain

$$P(\sigma_m, \sigma_{-m}) = \frac{2\sigma_m\theta + 2\theta(1 - \sigma_m) + \mu(1 - \sigma_m + \sigma_{-m}(1 - \sigma_m + \sigma_m))}{4\theta + 2\mu} = \frac{2\theta + (1 - \sigma_m + \sigma_{-m})\mu}{4\theta + 2\mu}$$

This probability is identical to the one displayed in Eq. (1.13). ∎

Proof of Lemma 1.3.   Since miners instantly sell their block rewards of value z_m and obtain linear utility from consuming the corresponding amount of numeraire good, we have that $W_M^m(z_m) = z_m + W_{0,M}^m$.

The component $W_{0,M}^m$ of the value function can be obtained by summing-up the expected payoffs obtained at each mining round, discounted at present value. Since miners consider the length of a mining round $T_{H+1}^* - T_H^*$ as exponentially distributed with rate μ, property (1.A.6) yields

$$\mathbb{E}\left[\psi \int_0^{T_{H+1}^* - T_H^*} e^{-rt}\, \mathrm{d}t \,\middle|\, T_H^*\right] = \psi \int_0^\infty e^{-(r+\mu)t}\, \mathrm{d}t = \frac{\psi}{r+\mu} \tag{1.B.5}$$

On the other hand, the present value of the block reward, discounted for the expected length of a mining round is obtained by computing

$$\mathbb{E}\left(e^{-r(T_{H+1}^* - T_H^*)} R(Q_{H+1})\right) \underbrace{=}_{g_t(Q_t)=g(Q_t)\forall t} \mathbb{E}\left(e^{-r(T_{H+1}^* - T_H^*)} R(Q)\right) \underbrace{=}_{\text{from (1.A.5)}} \frac{\mu}{r+\mu}\mathbb{E}\left[R(Q)\right] \tag{1.B.6}$$

with

$$\mathbb{E}\left[R(Q)\right] = g(0)\left[R(0)P_{M-1}(0,0)\right] + (1-g(0))\left[\sigma R(1)P(1,\sigma) + (1-\sigma)R(0)P_{M-1}(0,\sigma)\right]$$

$$= R(0)\left((1-\rho)P_{M'}(0,0) + \rho(1-\sigma)P_{M'}(0,\sigma)\right) + R(1)\sigma P_{M'}(1,\sigma) \tag{1.B.7}$$

The expression resulting by the combination of Eqs. (1.B.5) to (1.B.7) is

$$W_0^m = \frac{\mu\left(\mathbb{E}\left[R(Q)\right] + W_0^m\right) - \psi}{r+\mu}$$

$$\text{so that } W_0^m = \frac{\mu\mathbb{E}\left[R(Q)\right] - \psi}{r} \tag{1.B.8}$$

To obtain miners' value function (1.23), subtract from Eq. (1.B.8) the value of the initial investment in the mining node F and add real balances ∎

Proof of Lemma 1.4. Using property (1.A.7), we have

$$W^b(z) = -(z^* - z) + \frac{\mu W^b(z^*(1-\pi)) + \alpha V^b(z^*)}{r+\mu+\alpha}$$

$$(r+\mu+\alpha)W^b(z) = -(z^* - z)(r+\mu+\alpha) + \mu W^b(z^*(1-\pi)) + \alpha V^b(z^*)$$

$$rW^b(z) = -(z^* - z)(r+\mu+\alpha) + \mu\left[W^b\left(z^*(1-\pi)\right) - W^b(z)\right] + \alpha\left[V^b(z^*) - W^b(z)\right]$$

$$rW^b(z) = -(z^* - z)r + \mu\left[W^b\left(z^*(1-\pi)\right) - W^b(z^*)\right] + \alpha\left[V^b(z^*) - W^b(z^*)\right]$$

$$rW^b(z) = -(z^* - z)r - \mu\pi z^* + \alpha\left[V^b(z^*) - W^b(z^*)\right]$$

Expression (1.26) follows by setting $z = 0$ ∎

Proof of Lemma 1.5. Immediate from property (1.A.7) and the derivations in Appendix 1.A.1 ∎

1.C   Notation

t	time index	α	DM meeting rate
b, B; s, S; m, M	agents' index and population	x	generic good consumption
y	special good consumption	$u(\cdot)$	special good utility
$a_{i,t}$	nominal balances, $i \in \{b, s, m\}$	$z_{i,t}$	real balances, $i \in \{b, s, m\}$
B_{i}, S_{i}	subscribed buyers and sellers	$A_{i,t}^D$, $A_{i,t}^S$, $A_{i,t}$, A_t	token demand and supply
$Z_{i,t}$, Z_t	real balances in CM_i	$\phi_{i,t}$, ϕ_t	tokens' price in CM_i
$R(k)$	block reward for block size k	π	seigniorage rate
τ	proportional fee rate	h, H	block and blockchain height
$T_{h,m}$	publication time of block (h, m)	T_H^*	publication time of ref. block
μ, μ_m	block creation and mining PoW rate	$k_m(Q)$	block size Markov strategy
λ	transaction request rate	r	discount rate
$\underline{M}$	safe miners' participation	f, $\tilde{f}$	actual and approx. density
θ	block transmission rate	F, ψ	up-front and flow mining cost
$\mathcal{M}$	set of miners' labels	$\mathcal{U}_i$	life-time utility of agent
$\gamma_{h,m}$	PoW solution time	$\epsilon_{h,m}$	transmission lag
σ_m	prob. of filling a block	$p(y)$	price of the special good
Q	miners' mem-pools size	n	block number
$g(\cdot)$	p.m.f. of mempool size	$d(\cdot)$	p.m.f. of validation time
ρ	load factor	ν	probability of validation
β	buyers bargaining weight	Ω	validation discount rate

1.D   Glossary

Altcoins	Cryptocurrencies originated from forks of BTC
Bitcoin (BTC)	The original implementation of the Bitcoin blockchain
bitcoin	A Bitcoin token
blockchain	A copy of a distributed ledger
coinbase transaction	The transaction awarding a miner with new tokens
fork	A bifurcation of the blockchain
genesis block	The first block of a blockchain
hash	A code resulting from an encryption
hashpower	Time rate of hash codes produced by the mining node
height (block)	Distance between a generic block and the genesis block
longest chain rule (LCR)	The consensus formation rule employed by most PoW blockchains
mining node	A computer endowed with dedicated mining technology
mining pool	A consortium of miners
mining	The activity of recording blocks of transactions on the blockchain
mem-pool	Set of pending transactions received by a miner
proof-of-work (PoW)	A blockchain protocol based on computationally intensive record keeping
proof-of-stake (PoS)	A blockchain protocol based on a funds staking mechanism
satoshi	A Bitcoin token containing 10^{-8} bitcoins
stale block	A block lying in an abandoned chain
SHA-256	A common encryption algorithm

Page intentionally left blank.

Chapter 2

A Theory of Crowdfunding Dynamics

(Joint with Matthew Ellman)

2.1 Introduction

Empirical analyses of crowdfunding highlight a U-shape in the plot of average funding pledges against time (see Fig. 2.D.1). Understanding the dynamic interactions between funders that underlie these patterns will enable entrepreneurs to better exploit crowdfunding as a tool to fund their innovations. Yet first generation theories only model simultaneous funding choices. We build a parsimonious dynamic theory that can explain the U-shape. Our new ingredient is that funders must sink an inspection cost if they want to learn their value of the entrepreneur's innovation *before* it gets produced. This generates rich dynamics since bidders only pay to inspect when they perceive sufficient success prospects. In turn, inspection choices drive bidding which drives success. So we solve for the co-evolving rates of success and bidding. Our equilibrium characterization predicts how outcomes, including initial success rates and the slope of average bidding profiles, depend on project characteristics and design choices. Our framework identifies two novel dynamic effects.

The first effect is driven by a tendency for pivotality to decrease. Later bidders arrive nearer to the campaign deadline and therefore have fewer followers to influence. As a result, their average pivotality is lower than that of earlier bidders. This decreasing pivotality has a negative effect on the slope of average bidding because pivotality raises the benefit from inspecting which is a necessary precursor to bidding. The second effect is driven by the fact that later bidders learn more about other bidders' choices: they observe the sum of prior bids. High prior bidding represents good news for success so it increases later bidding while low bidding represents bad news and decreases later bidding. Whether the good news' positive effect dominates, or is dominated by, bad news' negative

effect depends, via Jensen's inequality, on the shape of the inspection cost distribution. We call this the Jensen effect of learning.

Under plausible conditions, the decreasing pivotality effect is strongest in early stages, whereas a dominant, positive Jensen effect is most common in late stages of a campaign. So these two insights alone can explain a U-shaped bidding profile. The argument hinges on the fact that campaign success prospects usually become increasingly clear over time. Success clarity implies that success rates approach zero or unity. In either case, pivotality stabilizes near its lower bound of zero with little scope for further decrease. Decreasing pivotality therefore generates an initial downward slope which later levels out creating the flat bottom of the U-shape. By contrast, the success clarity of the late stages generates a strong positive Jensen effect as maximal success rates activate reluctant bidders. This explains the final upward slope of the U-shape.

In a companion paper, we allow for endogenous timing of bidder moves and bidders who arrive before the campaign starts. This extension is interesting because it can explain sharp spikes at campaign start and expiration dates. In the discussion, we consider a number of other factors that further affect bidding patterns. None offer such a direct and parsimonious explanation of bidding slopes and the U-shape. We find it impressive that the parsimonious, exogenous move order model of the current paper can already generate a U-shape. The model also generates new testable predictions and allows us to tackle the challenge of optimal campaign design.

We focus on the classic case of *all-or-nothing reward*-based crowdfunding. An entrepreneur makes an open call to fund her plan to make a new product. She commits to *reward all* funders with her product if and only if aggregate funds reach a threshold by a deadline. That is the "success" event. Otherwise, the project fails and funds are returned; *nothing* is produced and nothing is paid. We treat the simplest case where there is one product and one price, so the entrepreneur proposes a single reward offer. Funders are essentially advance buyers whose bids or pledges are commitments to pay the product price *if* funds reach the threshold in time. Unlike competing bids in an auction, bids are strategic complements because a bid yields nothing without enough other bids to reach the threshold. Ellman and Hurkens (2019b) use mechanism design to prove that simultaneous bidding and the all-or-nothing paradigm are optimal when funders can be fully reimbursed on project failure. However, with sunk costs of inspection, sequentiality is needed for coordination to reduce wasteful inspections. Major crowdfunding platforms facilitate this by accepting bids over time and publicizing the cumulative aggregate.

To flesh out our two dynamic effects, we focus on an exogenous ordering of bidder moves. Bidders either cannot or do not want to delay decision-making. They "arrive" when they become aware of the entrepreneur's campaign. We assume a homogenous Poisson arrival process, so that the (average) bidding profile would be flat if there were no strategic interactions. Each bidder observes his arrival time and the bids accrued so

far. So he knows how much time is left till the deadline and the funding "gap" still needed to reach the threshold. The time and gap form the state of a Markovian sequential move game, because each bidder only cares about others' choices in so far as they affect the success probability when he bids. We call this bid-conditional probability, the bidder's success rate. It is the probability of success perceived by a bidder who arrives and bids in the current state. This rate plays a central role in the analysis. Sinking inspection costs is worthwhile precisely when this rate is high enough.[1] Notice that it equals the unconditional success rate plus the bidder's pivotality, the success impact of his bid.

By the law of iterated expectations, the unconditional success rate is always a martingale. So decreasing pivotality directly implies that the bidder's success rate is also decreasing. That in turn pushes towards a negative bidding slope. This is the decreasing pivotality effect (DPE). Formally, we establish decreasing pivotality by proving that pivotality is a supermartingale. The economic intuition, that later movers have *fewer* followers to positively influence before the campaign deadline hits, applies to any sequential move game with strategic complementarity and a fixed deadline. As noted above, the magnitude of this decreasing pivotality effect (DPE) usually falls over time as pivotality usually approaches its lower bound of zero in the last stages of a campaign. The cost distribution also matters by shifting followers' sensitivity to influence.

The overall expected effect of time combines this with the Jensen effect of learning about shocks to the gap. Uncertainties in arrivals, costs and tastes generate bidding shocks that lead to shocks in the gap and hence in pivotality and the bidder's success rate. Strategic complementarity magnifies the variance over time so that the bidder's success rate becomes increasingly dispersed. The net expected impact of these shocks is positive if the density of bidders is increasing in inspection costs and negative if the density is decreasing. This is intuitive: an increasing density, equivalent to convexity of the cumulative distribution function (CDF), implies increasing returns to good news. The Jensen effect then pushes the bidding slope upwards. Sufficient convexity can dominate the decreasing pivotality effect and yield a positive sloped bidding profile. Conversely, CDF concavity reduces expected inspection and hence bidding. Decreasing returns to good news push in the same direction as decreasing pivotality, so a linear or concave CDF guarantees a negative slope.

We derive a range of possible profile shapes. To explain the U-shape, we emphasize a context in which a majority of bidders have a continuous distribution of inspection costs with a density that is either decreasing in cost or takes any shape except one that is steeply increasing on a significant range of costs. The uniform distribution obviously satisfies the latter and a single-peaked density with a negative mode satisfies the former as well. For the remaining fraction of bidders, the effective inspection cost is an atom at the upper bound on the cost support that bidders are ever willing to pay. The actual cost

---

[1]We assume a price high enough that bidders never bid without first inspecting – *no "blind bidding"*.

could follow any distribution because the motivation is that this group of bidders follow a rule of thumb. They only inspect sure prospects and neglect the prospect to bid on a project for which feasibility depends on what other bidders do later. This setting naturally generates a U-shape. First, the standard distribution creates an initial downward slope that becomes flatter over time as pivotality tends to become less significant. Second, the effective limit cost types cause a Jensen effect that is positive and strong during the late phase of the campaign. This Jensen effect is negligible early on because it is rare for a campaign to reach a low enough gap early on in order to activate the rule of thumb types who effectively have the limit cost. This explains the final upward part of the U-shape.

To discuss further slope possibilities, note that costs below zero are equivalent to a zero cost since all such bidders inspect on arrival. Costs strictly above the expected inspection benefit given guaranteed success are equivalent to the bidder not arriving. So the relevant cost range is from zero to this upper cutoff. The Jensen effect is readily signed if the cost density is monotonic on this range. In the case of a single-peaked density, a negative mode implies that the peak cost is below zero and the density is decreasing on the relevant range, implying a concave CDF and a negative Jensen effect. Similarly, a mode above the cutoff guarantees a convex CDF and positive Jensen effect. Empirically, this suggests that the bid profile slopes downwards when inspection is an attractive distraction for most relevant bidders. Negative net inspection costs predominate when bidders' curiosity, active enjoyment in reading or desire for distraction tend to outweigh their costs of effort and opportunity cost of time. Also close contacts and family of the entrepreneur are likely to be anyway informed or to feel obliged to pay attention. Bidder connectedness with the entrepreneur, as estimated via Facebook links or geography (see Agrawal et al., 2011), can serve as a proxy for negative inspection costs.

The uniform distribution clearly implies a downward slope because the Jensen effect is always neutral and the decreasing pivotality effect (DPE) applies throughout. The intensity of DPE is non-monotonic in the initial gap. At low gaps, the project success rate rises and reaches unity so that pivotality falls to zero even for the earliest bidder. At very high gaps, the project success rate is very low so even the earliest, most pivotal bidders have minimal pivotality. Average pivotality still falls but it cannot fall very far since pivotality cannot fall below zero. Nonetheless, average pivotality and bidding plots are increasingly convex as the threshold rises.

Higher power distributions are sufficiently convex to produce a positive slope. The magnitude of the Jensen effect is largest at the start and rapidly declines because the greater pivotality of early bids implies a greater variance early on.

Discrete inspection cost distributions are instructive and easily simulated, but discontinuities in the CDF generate bidding discontinuities. The homogenous cost case is the simplest because all bidders make the same inspection choice at any given state. The project is either active with maximal bidding or frozen because no bidders inspect. Once

frozen, bidding falls to zero so the project cannot heat up again. The frozen state is absorbing so the probability that the project is frozen can only increase over time. This implies that bidding decreases over time since the average bidding rate at any date is proportional to the probability of survival until that date. The bidding decreases take the form of discrete drops at a finite set of critical dates at which the project would freeze for one of the gaps between the initial gap and two. Pivotality at the critical state given by a gap and its critical date is continuous in time but on average falls in the next instant so that average bidding suffers a discrete drop. The Jensen effect is positive but continuous and therefore dominated by DPE at critical dates. The Jensen effect is trivial in active non-critical states since all bidders continue to inspect at the next instant no matter what happens. The Jensen effect is also trivial in frozen states because freezing implies no further bid variation.

The homogenous case is very special for that reason: all bidding freezes up at the same instant. A binary distribution with one atom at zero cost is a natural extension away from homogeneity, given that entrepreneurs have friends, family, fans and contacts. In this case, we say that the project is in a hot state if both zero cost and positive cost types are inspecting, while it is cold if only the zero cost bidders inspect. Critically now, bid variation is non-zero in the cold state, so the Jensen effect after each critical date is non-trivial and it is positive because bid variation provides the chance that high cost bidders will start to inspect again . So average bidding rises back up gradually, only to drop down discretely at the next critical date. This generates a saw tooth pattern.

We can also apply our framework to answer key questions of optimal design. Our characterization of success rates reveals that the overall success rate at the start of a campaign is falling in the threshold and rising in its duration for a fixed arrival rate (though raising duration may, in principle, lower the arrival rate). The success rate is also decreasing in the crowdfunding price. Since we prove that a bidder's *pivotality* tends to be greater if he arrives earlier, it seems natural that an initially higher price discount should help the entrepreneur to promote project success. This is true for an entrepreneur without commitment, but when able to commit to future discounts, that is usually better still because future discounts motivate prior bidders as well as subsequent bidders. Unrestricted optimal pricing rules are more generally too complicated for practical use since the price would depend quite intricately on the bidding path but limited rewards are commonly used and we shed light on their optimal design.

We discuss related works in Section 2.1.1. Section 2.2 presents the baseline model with exogenous move orders. Section 2.3 characterizes success rates and pivotality and presents our main result on the slope of the average bidding profile. In Section 2.4, we apply the methodology to specific distributions of inspection costs: linear or uniform, quadratic and higher powers, homogenous costs and binary distributions. Section 2.5 presents our analysis of design, endogenizing parameters chosen by the entrepreneur.

Section 2.6 discusses realism and subtler issues regarding the modeling choices. We wrap up and discuss alternative and complementary theories in Section 2.7. Omitted proofs can be found in Appendix 2.A.

2.1.1   Related literature

Empirics.   Empirical studies on bid dynamics in crowdfunding are relatively recent as data was previously unavailable. Kuppuswamy and Bayus (2017) is the first paper to our knowledge to report that the U-shape (or bathtub) pattern of pledge revenue and count is pervasive across project categories and outcomes (i.e. success and fail) on Kickstarter (KS for brevity). Later Liu (2020) and Deb et al. (2021) corroborate that prior evidence. Crosetto and Regner (2018) instead find the U-shape on the German platform Startnext resulting from the aggregation of notable heterogeneity at campaign level. Deb et al. (2021) find that the spikes of the U-shape are amplified by an increased bidder arrival rate. In the starting phase, the increase is due to higher advertising and social activity, e.g. KS attributing the "Projects we Love" badge to campaigns it favours. In the conclusive phase, also bidders' strategic waiting plays an important role. Indeed, the final peak displayed in figure 1(c) of Deb et al. (2021) starts two days prior to campaigns' conclusion because that is the time at which the "Remind Me" option that Kickstarter offers to bidders reminds them to check back a chosen campaign that is running out of time. Deb et al. (2021) and Crosetto and Regner (2018) also break-down the pledge U-shape into purchase and donation income and find that purchases prevail early for early-finisher campaigns while donations dominate late for later-finishers. Deb et al. (2021) highlight the empirical puzzle of a pledge revenue drop occurring right after campaigns reach their goal and attribute its possible causes to a successive price increase or reduction in advertising and social activity.

Success rate prediction based on project static and dynamic features is a salient empirical challenge. The overall picture is that the success rate of KS campaigns is around 40%. in predicting success Rao et al. (2014) and Etter et al. (2013) predict success using KS data. The first group of authors uses bidding rates and their slope as predictors. They construct strong success predictors using the bid rate computed over the first 10% and between 40 and 60% of the campaign's duration and the bid slope over the final 5%. Etter et al. (2013) predicts success with 76% accuracy after the first four hours combining bid and social activity data. Social factors account for a 4% increase in prediction performance. Both authors predict success with 84 and 85% accuracy using bid revenue over the first %15 of the campaign's duration. Crosetto and Regner (2018) find that initially failing campaigns are often saved at a late stage by intensified communication and support from the entrepreneur's extended personal network. Similarly, Colombo et al. (2015) associates intense initial bidding with social capital developed within the crowdfunding

community.

A good number of studies analyze static (project level) features of crowdfunding campaigns. Agrawal et al. (2014) and Kuppuswamy and Bayus (2018) provide a survey. Agrawal et al. (2011) show that funders who are geographically close to the entrepreneur tend to bid early on Kickstarter. Mollick (2014) and Cordova et al. (2015) highlight the features of successful crowdfunding practices. Mollick (2014) uses KS data. Cordova et al. (2015) form a data-set by combining data on technology projects active in 2012 from Kickstarter, Indiegogo, Eppela and Ulule. Both teams find that higher goals are associated with lower success rates, but are discordant on the success impact of campaign duration, which is positive according to Cordova et al. (2015) but negative for Mollick (2014). Both teams observe bimodality in the distribution of total pledge income and completion time.

Theory. Early crowdfunding models (Belleflamme et al., 2014; Chang, 2016; Chemla and Tinn, 2018; Ellman and Hurkens, 2019a,b; Hu et al., 2015; Sahm, 2016; Strausz, 2017) are static apart from a post-crowdfunding sales period. Ellman and Hurkens (2019b) demonstrate crowdfunding's role in credibly learning market demand and use mechanism design to prove optimality in the case with two possible valuations, high and low. Strausz (2017) restricts to settings where the low valuation equals zero to study optimal design when the entrepreneur can abscond with the funds instead of producing; see also Chang (2016); Chemla and Tinn (2018); Ellman and Hurkens (2019a). We abstract from moral hazard and similarly restrict to the case where the low valuation is zero to focus on the dynamic challenge.

A more recent literature explicitly models pledging as a dynamic subscription game. In this context inspection costs are natural because funders bid before the innovative good has been produced. This provides a plausible justification for some contemporaneous models of crowdfunding dynamics Alaei et al. (2016); Deb et al. (2021). Those other papers assume an exogenous and deterministic bidding cost, motivated as an opportunity cost of funds, which is often small.

Alaei et al. (2016) study optimal crowdfunding design and the dynamics of bids assuming bidders arrive exogenously with binary valuations, facing a homogeneous opportunity cost to bid. We treat the analogous continuous-time problem, adding cost heterogeneity. They do not derive the average bid profile, but our analysis proves that it is necessarily decreasing in their homogenous cost setting.

Deb et al. (2021) seek to explain a U-shaped bidding profile by adding a distinct type of bidder: buyers arrive and bid exactly as in Alaei et al. (2016) and our model, but there is also a (single) donor who is always present and can donate any amount at any time (the donor bids in every subperiod after a buyer arrives or fails to arrive). As in Alaei et al. (2016), buyers are homogeneous and face an opportunity cost of pledging. Bid variation is driven by buyer uncertainty about the donor's wealth as well as the randomness in

buyer arrivals. The donor is vital for explaining the U-shaped dynamics in their paper. Our analysis proves that the profile is otherwise decreasing for their setup, as just noted in the comments on Alaei et al. (2016). By studying cost heterogeneity, we derive rich dynamics without need for any donor types. We abstract from the complex signalling game associated with donor wealth uncertainty and the power to bid continuously over time. We are thereby able to analyze the shape of average bidding profiles. Deb et al. (2021) instead focus on when bidding can feature spikes, defined as multiple bids made in one instant. By assumption, bidders cannot be responsible for such an occurrence because bidders only ever make one bid and the probability that two bidders arrive at an identical date in the continuous time limit is zero. The donor can make multiple bids and the analysis proves that the donor makes just one bid at a time except at the very start and very end of a campaign.

Common values offer another promising avenue for understanding dynamics. Given the prominent role of thresholds and limited campaign duration in actual crowdfunding, it will be important to combine common values with the model that we develop below. Nonetheless, there are already some suggestive early contributions. Liu (2020) endogenizes move order, but only in a two-period game. Assuming common value, she shows that equilibria in cutoff strategies with respect to private signals (coordinated equilibria), predict that successful ventures are generally more back-loaded than failed campaigns. She notices that the game does not have a unique equilibrium prediction but the equilibria she analyses address the research question of what drives selection into leaders and followers in a collective action. Babich et al. (2017); Vismara (2016); Zhang and Liu (2012) provide valuable related contributions in the context of equity-based crowdfunding.

Other papers model funding dynamics but concentrate on crowdfunding design. Zhang et al. (2017) and Chakraborty and Swinney (2019) consider a bidder crowd divided into ordinary and herder bidders. These last wait and make a last-minute action. In Zhang et al. (2017) types are exogenous while in Chakraborty and Swinney (2019) bidders choose strategically to wait but are partially rational. Their model predicts a final bidding spike but cannot explain a decreasing profile.

All these closest contemporaneous studies except Zhang et al. (2017) assume an exogenous sunk cost of making a bid. The main interpretation is that bids are held in escrow so bidders must forego other opportunities for using the money. This argument has some disadvantages. For the cost to matter it must be sunk even in the failure event where bids are reimbursed. So the opportunity cost is driven by fleeting substitute opportunities that arise during the campaign and disappear before it ends. The risk of such costs on a typical campaign lasting 40 days are generally small, especially when buyers are not credit constrained. Foregone interest earning opportunities also imply only a small cost. Such costs would be increasing in the relevant time till the deadline, $\tau - t$, contrary to the simple case of a fixed cost on which all the other studies have focused.

Our inspection cost perspective avoids these problems and provides an added reason to expect substantial variation in costs across bidders. We also complement those studies because inspection costs deliver qualitatively similar results. Inspections costs provide a more plausible motivation than the direct bidding costs assumed in all the related papers. Moreover, inspection costs justify assuming a time-independent distribution. In this way, our study complements those mentioned above.

Finally, there are several papers on donation-based crowdfunding. Solomon et al. (2015) study the timing of pledges experimentally. Donors tend to bunch their pledges at the beginning and at the end of a campaign. In line with our findings on endogenous timing, early pledging is a better strategy for donors who value success. Cason and Zubrickas (2018) model how refund bonuses affect public good provision and test their theory in a laboratory experiment. Bidders can pledge more than once in their continuous-time dynamic contribution game. They find that refund bonuses reduce momentum but do not predict the benchmark bid profile.

2.2 Model

We study the simplest reward-based crowdfunding campaign in order to study bid dynamics in continuous time. The entrepreneur specifies a product as the *reward* for crowdfunders, a price p of that single reward, a funding threshold or goal and a deadline of her campaign. Funders can pledge p if they hear about the campaign before its deadline. The campaign is said to *succeed* if and only if the sum of these pledged funds (restricted to p or 0) reach or exceed the threshold. In the success event denoted $\mathcal{S}$, each funder pays his pledge p to the entrepreneur who must invest in production and deliver the unit reward to each funder. In the complementary event of a *failure*, denoted $\mathcal{F}$, there is no production and no payments. We assume perfect enforcement of these crowdfunding rules, both on the entrepreneur and on the funders. Thus a pledge is a binding commitment to pay p in return for the product *contingent* on funding success $\mathcal{S}$. Admati and Perry (1991) call the pledges, subscriptions but we use the briefer word *bids*.[2]

Setting time t to zero at campaign initiation, the campaign's deadline equals its duration, denoted τ. Since funders bid p or nothing, a goal of pg_0, or any goal $p\tilde{g}_0$ with $\tilde{g}_0 \in (g_0 - 1, g_0]$, is equivalent to a threshold g_0 on the number of bids. This g_0 is also the initial *gap* between the threshold and cumulative bids since bids collected equals zero at the start. The campaign follows the predominant *All-or-Nothing* paradigm: in event $\mathcal{S}$, the campaign succeeds in collecting g_0 bids by τ, in which case all bids are paid, production occurs and all bidders receive one unit of the product as reward; otherwise, the campaign fails, event $\mathcal{F}$, so there is no production, no rewards and no

---

[2]In crowdfunding with unlimited units, bids are complements; they do not compete as in an auction.

payments (any bids held in escrow are reimbursed). The entrepreneur's production costs and possible credit needs affect how she chooses price p, duration τ and threshold g_0, but until design Section 2.5, we analyze dynamics within the bidding phase for exogenous project variables.

We say that a bidder arrives at the project at date t if he hears about the project at some $t \leq \tau$. Bidders arrive at a constant Poisson intensity $\lambda > 0$ throughout the campaign, giving an expected total of $\lambda\tau$ bidders. Neglecting zero probability events where multiple bidders arrive at an identical instant $t \in [0, \tau]$, we uniquely associate each bidder with his arrival date t. On arrival, each bidder perceives project characteristics τ, g_0, p, his arrival time t and the Poisson arrival intensity λ, but is uncertain about his private valuation v_t of the product. He can learn v_t immediately if he pays his inspection cost c_t, which he perceives on arrival. c_t represents t's cost of reading about the proposed product and introspecting.

Bidders immediately perceive on arrival the following common and fully independent distributions of these costs and values. For each t, $v_t \in \{0, v\}$ with probability $q \in (0, 1)$ on v and c_t has cumulative distribution function CDF, $F(\cdot)$.[3] Until we endogenize pricing, we normalize the price discount $d \triangleq v - p$ to unity; that is, we set $p = v - 1$ so that $d = 1$. Given this net benefit from buying the valued good of $d = 1$, a bidder's expected net benefit from buying contingent on learning $v_t = v$ is q. This justifies restricting the support of c_t to $[0, q]$: (a) negative costs, as from curiosity, enjoyment or caring for the entrepreneur, are equivalent to a zero cost, since we assume bidders inspect when indifferent; (b) a bidder with cost strictly above q is equivalent to a non-arrival, so any support above q is equivalent to reducing arrival rate λ.

Bidders are risk neutral. Each t gets 0 by not bidding, $v_t - p$ by bidding p if the project succeeds (event $\mathcal{S}$) and 0 by bidding on a project that fails. So bidding without inspecting gives him an expected payoff conditional on $\mathcal{S}$ of $qv - p$ and 0 if $\mathcal{F}$; there is no resale. We assume $p > qv$ to focus on the plausible case where bidders never blindly bid without inspecting to check that they value the good. As $p = v - 1$, this is equivalent to

Assumption 2.1 (*No blind bidding, NBB*). $q < 1 - 1/v$

Bidders observe a single summary of prior bidder behaviour: on arrival, he observes how many bids B_t have been collected before t. So he knows the funding gap $g_t \triangleq g_0 - B_t$. The entrepreneur sets g_0 but g_t evolves over time as a function of bidder arrivals and choices. We work with this gap, instead of the equivalent bid count B_t, so the project's publicly observable state at t is (t, g_t).

Move orders are exogenously determined by arrival dates t. Each bidder takes all his decisions in a single episode of negligible duration so bidders move sequentially. On arrival, a bidder either: (**A**) **A**voids the project and **A**voids bidding, (**B**) **B**lind bids in

---

[3]A non-zero low valuation, if still below p, leads to identical outcomes.

that he bids without inspecting to see v_t, (**C**) **C**hecks out the project by paying c_t and bids if he learns that $v_t = v$. Inspecting but bidding irrespective of v_t or bidding only when $v_t = 0$ are strictly dominated so we need only consider **A**, **B**, **C** which reduce to **A**, **C** under Assumption 2.1.

Remark. This assumes that bidders either cannot revisit the project later, prefer to decide in one go or react impulsively without thinking whether to delay their choice. When a project captures the bidder's attention at t, he registers its salient characteristics and it is often more efficient for him to decide how to react while his mind is already focused within a single thinking episode. Otherwise, he must decide to check back, later remember to do so and re-focus to finally decide. Curiosity and impulsiveness also push towards deciding in one go. In addition to being realistic for many bidders, exogeneity with type independence ensures that a project's success prospects at date t only depend on its current state (t, g_t), not the full history of prior arrivals and their types and inspection and bidding choices. In particular, equilibria would not change if bidders observed the full history of gaps up to t instead of only the current gap g_t. On the other hand, some platforms offer a "Remind Me" option for people who feel attracted to a project but are too busy or doubtful and prefer to wait to see how others bid before possibly inspecting further. Bidders using this option still need to engage in at least two thinking episodes but the informational benefit may compensate them. Analysis of settings with endogenous delays is much more complicated so we treat that case in a model variant called ENDO in Ellman and Fabi (2021a).

In sum, bidders arrive sequentially at Poisson arrival times t; each bidder t observes $\tau, \lambda, q, F(\cdot)$ and c_t and state (t, g_t) on arrival and chooses between substrategies **A** and **C**. The currently superfluous notation B_t and v, p, d reappear in the design analysis.

Equilibrium concept. We solve for undominated Perfect Bayesian Equilibria, abbreviated to PBE. Undominated refers to the fact that we restrict to weakly undominated strategies. For concreteness, we tie-break in favour of **C** among equilibria that generate identical payoff distributions. Thanks to the sequentiality of moves, this leads to a unique PBE. So there is no strategic uncertainty, but the three exogenous sources, Poisson arrivals, private inspection costs and valuations create the aggregate demand uncertainty underlying the shocks to success prospect S that are central to our analysis. Of these three sources, valuation uncertainty is fundamental. Without it, inspection would be pointless and without inspection costs, the project state (t, g_t) would not affect the choice between **A**, **B** and **C**.[4] Uncertain inspection costs create richer dynamic possibilities and are crucial for a positive slope. Arrival uncertainty serves to keep arrivals finite in a continuous time setting.

---

[4]Exogenous bidding costs c_t of **B** is equivalent to **C** with $q = 1$, neglecting Assumption 2.1.

2.3   Analysis

To solve for bidding equilibria given any campaign g_0, τ and $p = v - 1$ with parameters $\lambda, q, F(\cdot)$, we study the incentives of a bidder t after he observes his inspection cost c_t and campaign state (t, g), indicating that $g_t = g$. State (t, g) matters to bidders purely via its impact on the project's success prospects and bidders only affect each other via this success rate.[5] Using t_+ to indicate infinitesimally after t, that subgame starts at (t_+, g) if he does not bid and at $(t_+, g - 1)$ if he does bid. We define the evolving, state-contingent probability of success

$$S_{(t,g)} \triangleq \mathbb{P}(\mathcal{S}|(t,g)) \equiv \mathbb{E}_{(t,g)} \left(\mathbb{1}_{g_{\tau_+} \leq 0} \right) \tag{2.1}$$

where $\mathbb{E}_{(t,g)}(\cdot) \triangleq \mathbb{E}(\cdot|(t,g))$. This conditions on knowing $g_t = g$ but not whether a bidder arrives at t. If there *is* a bidder at t, the success probability rises to $S_{(t,g-1)}$ if he bids and stays at $S_{(t,g)}$ if not.[6] In choosing what to do, he only cares about success prospects if he bids. We denote this bid-contingent success probability by

$$S_{(t,g)}^{\mathrm{bid}} \triangleq S_{(t,g-1)} \tag{2.2}$$

The *pivotality* of a bidder arriving at t is $\Delta S_{(t,g_t)}$ where difference operator Δ denotes the impact of a unit <u>reduction</u> in g; $\Delta Y_{(t,g)} \triangleq Y_{(t,g-1)} - Y_{(t,g)}$ for a generic function Y. So

$$S_{(t,g)}^{\mathrm{bid}} \equiv S_{(t,g)} + \Delta S_{(t,g)} \tag{2.3}$$

This decomposition is helpful because $S_{(t,g_t)}$ is a martingale: by the Law of Iterated Expectations, $\mathbb{E}_{(t,g)} \left(\mathbb{E}_{(x,g_x)} \left(\mathbb{1}_{g_{\tau_+} \leq 0} \right) \right) = \mathbb{E}_{(t,g)} \left(\mathbb{1}_{g_{\tau_+} \leq 0} \right)$ for any $x \in [t, \tau]$, so, by Eq. (2.1),

$$\mathbb{E}_{(t,g)}(S_{(x,g_x)}) = S_{(t,g)} \tag{2.4}$$

Section 2.3.3 will prove that pivotality $\Delta S_{(t,g_t)}$ is a supermartingale which drives towards decreasing bidding profiles. First, we show how $S_{(t,g_t)}^{\mathrm{bid}}$ determines bidder t's choice.

Bidder t's simplest option is to choose to **A**void the project entirely by doing nothing, **A**. Not inspecting and not bidding always gives the same payoff, $u_t^{\mathbf{A}} = 0$. So for any (t, g), t's expected utility from **A** is $U_{(t,g)}^{\mathbf{A}} = 0$.

The **B**lind bid choice **B**, of bidding without inspecting, carries the risk of paying $p = v - 1$ despite not valuing the reward. So $u_t^{\mathbf{B}} = \mathbb{1}_{g_{\tau_+} \leq 0} (v_t - (v - 1))$. Taking expectations

---

[5]By type independence, bidders infer nothing about their own or future bidder valuations and costs. Bidder t takes the strategies of later bidders as given by the PBE; they cannot detect deviations.

[6]We can write t instead of t_+ because $S_{(t,g)}$ is continuous in t for any gap g.

and applying Assumption 2.1 proves that **A** dominates **B** as claimed above since

$$U^{\mathbf{B}}_{(t,g)} \triangleq \mathbb{E}_{(t,g)}\left(u^{\mathbf{B}}_t\right) = S^{\text{bid}}_{(t,g)}\left(qv - (v-1)\right) < 0$$

Bidding is the only weakly undominated action of a bidder t who learns that $v_t = v$. Conversely, not bidding is his only undominated action if he knows $v_t = 0$. So the only relevant substrategy involving inspection is **C**: **C**heck by paying c_t and bid if and only if the inspection reveals that $v_t = v$. This yields ex post payoff and expected utility,

$$u^{\mathbf{C}} = \mathbb{1}_{g_{\tau_+} \leq 0}\left(v_t - (v-1)\right)^+ - c_t$$
$$U^{\mathbf{C}}_{(t,g)} = qS^{\text{bid}}_{(t,g)} - c_t \tag{2.5}$$

In sum, a generic bidder's strategy is a mapping from each possible observed history or time, gap and cost trio to this trio of choices, $a : (t,g,c) \mapsto \{\mathbf{A}, \mathbf{B}, \mathbf{C}\}$. For a PBE, the choice a must be a best response at every possible information set (t,g,c) and bidder beliefs must be consistent with Bayes rationality at every information set.[7] Bayes rationality simply requires that each bidder's belief about the success probability given he bids is the correct state-contingent probability assessment S^{bid} derived below.

As **A** dominates **B**, we need only compare **A** and **C**. Eq. (2.5) shows that **C** is chosen whenever $c_t \leq qS^{\text{bid}}_{(t,g)}$. This has probability $F(qS^{\text{bid}}_{(t,g)})$. Since bidders arrive with Poisson intensity λ and **C** results in a bid with probability q, this generates non-homogenous Poisson bidding intensity,

$$\beta_{(t,g)} \triangleq \lambda q F(qS^{\text{bid}}_{(t,g)}) \tag{2.6}$$

Arrival rate λ and taste parameter q are fixed, so the systematic variations in $S^{\text{bid}}_{(t,g_t)}$ and resulting inspection rate $F(qS^{\text{bid}}_{(t,g_t)})$ fully determine the temporal pattern of bidding. We now analyze these co-moving variables.

2.3.1 The co-evolution of success probabilities and bids

To characterize how the bid-conditional success probability $S^{\text{bid}}_{(t,g)} \equiv S_{(t,g-1)}$ and bid intensity $\beta_{(t,g)}$ depend on parameters and the state, we study their interaction and solve for $S_{(t,g)}$ for any $g \in \mathbb{Z}$ and $t \in [0, \tau]$.

By the definition of a success, $S_{(\tau,g)} = 1$ for $g \leq 0$. Given that g can only decrease over time, $S_{(t,g)} = 1$ for $g \leq 0$ and any $t \leq \tau$. From this initial condition, we solve for higher gaps via a recursive step grounded in two facts. First, the bidding rate $\beta_{(t,g)}$ at gap g depends on $S^{\text{bid}}_{(t,g)}$ and hence on the success rate at gap $g - 1$. This is clear from Eqs. (2.2) and (2.6). Second, all paths leading to a success from any (t,g) with $g \geq 1$ must pass through the state $(T, g - 1)$ where $T \in (t, \tau]$ is the stopping time at which

---

[7]All the sets are reached with positive probability except in trivial cases where success is impossible.

the next bid arrives.[8] The ensuing success probability from that state is $S_{(T,g-1)}$. This reduced gap $g - 1$ also applies to the stopping time density as we now prove.

The stopping time of the next bid to arrive after reaching state (t, g) is T if there is no bid on $[t, T)$ and a bid at T. The gap stays at g over this time period, dropping to $g-1$ at the last instant T. By our convention, $g_T = g$ and $g_{T_+} = g - 1$. So T's density equals the bidding intensity $\beta_{(T,g)}$ at (T, g) times the probability $n_T^{(t,g)}$ of no bid on interval (t, T),

$$n_T^{(t,g)} \triangleq \exp\left(-\int_t^T \beta_{(x,g)}\, \mathrm{d}x\right) \tag{2.7}$$

The success rate from $(T, g - 1)$ is $S_{(T,g-1)}$. So, taking expectations over T given (t, g),

$$S_{(t,g)} = \int_t^\tau n_T^{(t,g)} \beta_{(T,g)} S_{(T,g-1)}\, \mathrm{d}T \tag{2.8}$$

This yields a unique solution for $S_{(t,g)}$ given a unique solution for $S_{(t,g-1)}$. This recursive derivation provides a unique solution for any g because the success rate is identically equal to unity and therefore unique for any $g \leq 0$. Together with Eq. (2.6), this proves

Proposition 2.1. *The crowdfunding game has a unique PBE characterized by* $a_{(t,g,c)} = \boldsymbol{C}$ *if and only if* $c \leq q S_{(t,g)}^{\mathrm{bid}} \triangleq q S_{(t,g-1)}$, *generating bid intensity* $\beta_{(t,g)}$, *where*

$$\beta_{(t,g)} = \lambda q F\left(q S_{(t,g)}^{\mathrm{bid}}\right) \equiv \lambda q F\left(q S_{(t,g-1)}\right) \tag{2.9}$$

$$S_{(t,g)} = 1 \text{ for } g \leq 0 \tag{2.10}$$

$$S_{(t,g)} = \int_t^\tau \exp\left(-\int_t^T \beta_{(x,g)}\, \mathrm{d}x\right) \beta_{(T,g)} S_{(T,g-1)}\, \mathrm{d}T \text{ for } g \geq 1 \tag{REC-S}$$

Notice that the recursion, like the bidders, only depends on the current states and beliefs about future bidding up till τ, so $S_{(t,g)}(\tau)$ is invariant to g_0 and changes in t and τ that fix $\tau - t$; $S_{(t,g)}(\tau) \equiv S_{(0,g)}(\tau - t)$. We work with time instead of time remaining since time moving forward is more intuitive. This proposition shows how the success rate affects bidding and how bid rates in turn affect success but the recursive formulation is relatively opaque. To better understand bidding dynamics, the next subsection uses the martingale property to relate success rate dynamics to bidding and pivotality.

2.3.2   Basic dynamic properties of the success rate

The fact that success rate $S_{(t,g_t)}$ is a martingale by Eq. (2.4) allows us to derive a differential equation linking success rate dynamics to bid rates and pivotality. We can also infer that pivotality is weakly positive while the partial time derivatives of unconditional and bidder's success rates are both negative.

---

[8]Later, we explicitly distinguish stopping times by bid number; here $T = T_{1+g_0-g}$.

At any instant, there are essentially two possibilities: either a new bid is collected, reflected in a unit drop in the gap from $g_t = g$ to $g_{t_+} = g - 1$, or the gap stays fixed. A unit drop requires a bidder to arrive, inspect and learn that he values the product. Formally, over an infinitesimal interval $[t, t + dt)$, neglecting order $(dt)^2$ terms, one bid arrives with probability $\beta_{(t,g_t)} dt$ or no bid is collected. Since $S_{(t,g_t)}$ is a martingale, $S_{(t,g)} = (1 - \beta_{(t,g)} dt)S_{(t+dt,g)} + \beta_{(t,g)} dt \, S_{(t+dt,g-1)}$. With $\dot{S}_{(t,g)}$ denoting the partial time derivative, this yields the ordinary differential equation (ODE) for S,

$$\dot{S}_{(t,g)} \triangleq \frac{\partial S_{(t,g)}}{\partial t} = -\beta_{(t,g)} \left(S_{(t,g-1)} - S_{(t,g)} \right) \qquad \text{(ODE-}S\text{)}$$

That is, $\dot{S}_{(t,g)} \equiv -\beta_{(t,g)} \Delta S_{(t,g)}$.[9] The jump term for $S_{(t,g_t)}$ is precisely the *pivotality* $\Delta S_{(t,g_t)}$ of a bid at t. So Eq. (ODE-S) captures the fact that the news effect or pivotality of a bid, weighted by bid intensity, exactly counterbalances the effect of time passing on success. We can sign both these effects.

Lemma 2.1 (Success rate properties)**.** *(i)* $\Delta S_{(t,g)}, \Delta S_{(t,g)}^{bid} \geq 0$, *(ii)* $\dot{S}_{(t,g)}, \dot{S}_{(t,g)}^{bid} \leq 0$.

The proof in Appendix 2.A uses induction on g. Intuitively, pivotality is never negative: each additional bid is good news for success, weakly raising S by reducing the remaining funding gap. Conversely, no news is bad news: time passing with no new bids lowers S since it leaves less time to cover the current gap by the deadline τ. Together with Eq. (2.6), Lemma 2.1 proves that higher gaps lower success and bidding rates:

Corollary 2.1 (The Gap)**.** *For any t, $S_{(t,g)}$, $S_{(t,g)}^{bid}$ and $\beta_{(t,g)}$ are all weakly decreasing in g.*

Three steps remain for drawing conclusions about time profiles. We need to take account of how g_t changes with time. Lemma 2.1 shows that $S_{(t,g)}^{bid}$ falls with t for a fixed gap g but since g_t falls over time, the overall average impact is not obvious. Nonetheless, in Section 2.3.3, we derive an unambiguous effect. Then in Section 2.3.5, we derive implications for $\beta_{(t,g_t)}$ via Eq. (2.6) and Jensen's inequality. Finally, in Section 2.3.6, we compute g_t distributions to convert results on instantaneous slopes from a generic state (t, g) into slopes of the average bidding profile.

2.3.3 Decreasing pivotality

Given that $S_{(t,g_t)}$ is a martingale, any systematic effect of time on $S_{(t,g_t)}^{bid}$ must come from the pivotality term $\Delta S_{(t,g_t)}$. We now prove that expected pivotality decreases over time. The logic behind decreasing pivotality derives from strategic complementarity and the fact that later arrivals have fewer bidders to influence. First, a bid at t strategically

---

[9]This also follows by differentiating Eq. (REC-S). Conversely, Appendix 2.C derives Eq. (REC-S) from Eq. (ODE-S) using n_T of Eq. (2.7) as integrating factor.

complements inspection and bidding by *subsequent* bidders. Second, an early bidder, low t, has more such followers to influence. This generates a larger average impact on success. This Decreasing Pivotality has a negative Effect (the DPE) on the slope of the bid rate as we will show in Section 2.3.5.

Formally, the change in pivotality at future date x expected from initial state (t, g),

$$\mathcal{D}_x^{(t,g)} \triangleq \mathbb{E}_{(t,g)}(\Delta S_{(x,g_x)}) - \Delta S_{(t,g)} \tag{2.11}$$

is always negative or zero. We prove this by showing that $\Delta S_{(t,g_t)}$ is a supermartingale via the infinitesimal generator for studying expected rates of change of a generic stochastic process $Y_{(t,g_t)}$: $\mathcal{L}_{(t,g)}^Y \triangleq \lim_{dt\downarrow 0} \left(\frac{1}{dt} \left(\mathbb{E}_{(t,g)} \left(Y_{(t+dt,g_{t+dt})} \right) - Y_{(t,g)} \right) \right)$. In the case of pivotality,

$$\mathcal{D}^{(t,g)} \triangleq \lim_{dt\downarrow 0} \left(\frac{1}{dt} \left(\mathcal{D}_{t+dt}^{(t,g)} \right) \right) \equiv \mathcal{L}_{(t,g)}^{\Delta S} \tag{2.12}$$

The key result on generators from Itô's formula (mathematical details in Appendix 2.C),

$$\mathcal{L}_{(t,g)}^Y = \dot{Y}_{(t,g)} + \beta_{(t,g)} \Delta Y_{(t,g)} \tag{GEN}$$

is very intuitive: $\dot{Y}_{(t,g)}$ captures time's direct effect while $\Delta Y_{(t,g)}$ weighted by the bidding intensity $\beta_{(t,g)}$ captures time's expected impact via negative unit jumps in the gap g_t.

Appendix 2.C also shows (via Dynkin's theorem) that $Y_{(t,g_t)}$ is a martingale if and only if its generator is identically zero, i.e., $\mathcal{L}_{(t,g)}^Y \equiv 0, \forall t, g$. Similarly, supermartingales and submartingales respectively correspond to everywhere weak negativity and positivity of $\mathcal{L}_{(t,g)}^Y$. Note that $\mathcal{L}_{(t,g_t)}^S = 0$ as $S_{(t,g_t)}$ is a martingale, so (GEN) immediately reconfirms Eq. (ODE-S). We now apply (GEN) to pivotality $\Delta S_{(t,g_t)}$ to prove that $\mathcal{D}^{(t,g)} \leq 0$.

Proposition 2.2 (Decreasing pivotality). *$\Delta S_{(t,g_t)}$ and $S_{(t,g_t)}^{\mathrm{bid}}$ are supermartingales.*

Proof. As $S_{(t,g_t)}$ is a martingale, Eq. (2.3) implies $\mathcal{L}_{(t,g)}^{\Delta S} = \mathcal{L}_{(t,g)}^{S^{\mathrm{bid}}}$. By (ODE-$S$)$|_{g-1}$, $\dot{S}_{(t,g-1)} = -\beta_{(t,g-1)} \Delta S_{(t,g)}^{\mathrm{bid}}$. So applying (GEN) to $S_{(t,g_t)}^{\mathrm{bid}}$,

$$\mathcal{D}^{(t,g)} = \mathcal{L}_{(t,g)}^{\Delta S} = \mathcal{L}_{(t,g)}^{S^{\mathrm{bid}}} = \dot{S}_{(t,g-1)} + \beta_{(t,g)} \Delta S_{(t,g)}^{\mathrm{bid}} = -\Delta\beta_{(t,g)} \Delta S_{(t,g)}^{\mathrm{bid}} \leq 0 \quad \blacksquare \tag{2.13}$$

The formula $\mathcal{D}^{(t,g)} = -\Delta S_{(t,g)}^{\mathrm{bid}} \Delta\beta_{(t,g)}$ neatly captures the insight that a later bidder has fewer successor bidders to encourage with his bid. For further intuition, we compare a bidder who bids at (t, g) and a bidder who bids at $(t + dt, g)$. The earlier bidder lowers the gap to $g - 1$ at t instead of at $t + dt$ and therefore differentially encourages any successor who arrives during that marginal interval. The probability of this interim successor arriving is $\lambda\, dt$. This successor is encouraged to bid with increased probability $q\Delta F(qS_{(t,g)}^{\mathrm{bid}}) = \Delta\beta_{(t,g)}/\lambda$. The success impact of that extra bid is $\Delta S_{(t,g-1)} = \Delta S_{(t,g)}^{\mathrm{bid}}$ since the gap fell to $g - 1$ at t. So the marginally earlier bidder's greater expected pivotality is

$\Delta\beta_{(t,g)}\Delta S^{\mathrm{bid}}_{(t,g)}\,\mathrm{d}t$ and the rate of expected change $\mathcal{D}^{(t,g)} = -\Delta S^{\mathrm{bid}}_{(t,g)}\Delta\beta_{(t,g)}$.

This formula is useful in Section 2.4. The systematic decreasing pivotality pushes via Eq. (2.6) towards a negative sloped bidding profile in Section 2.3.4. Since uncertainty will play a key role in Section 2.3.5, we end with an illustration (which the reader can safely skip) of the substantial stochastic variation behind this pivotality trend. In particular, pivotality *can* attain its maximal and minimal values in the last moments of a campaign.

> **Illustration.** It is perhaps surprising that average pivotality decreases from *any* state (t, g). If g_t does not change, $S^{\mathrm{bid}}_{(t,g_t)}$ falls with t by Lemma 2.1 but $S^{\mathrm{bid}}_{(t,g_t)}$ rises whenever g_t falls by Corollary 2.1. What Proposition 2.2 shows is that the average direct effect of time always dominates the positive average indirect effect of downward jumps in the gap and does so for any cost distribution. Fig. 2.1 probes this averaging effect for a project with initial gap $g_0 = 10$, duration $\tau = 20$, bidder arrival intensity 0.95, valuation probability $q = 0.8$ and uniform cost distribution on $[0, q]$. $S^{\mathrm{bid}}_{(t,g)} \equiv S_{(t,g-1)}$ and pivotality $\Delta S_{(t,g)} \equiv S_{(t,g-1)} - S_{(t,g)}$, so we only plot $S_{(t,g)}$ against time.
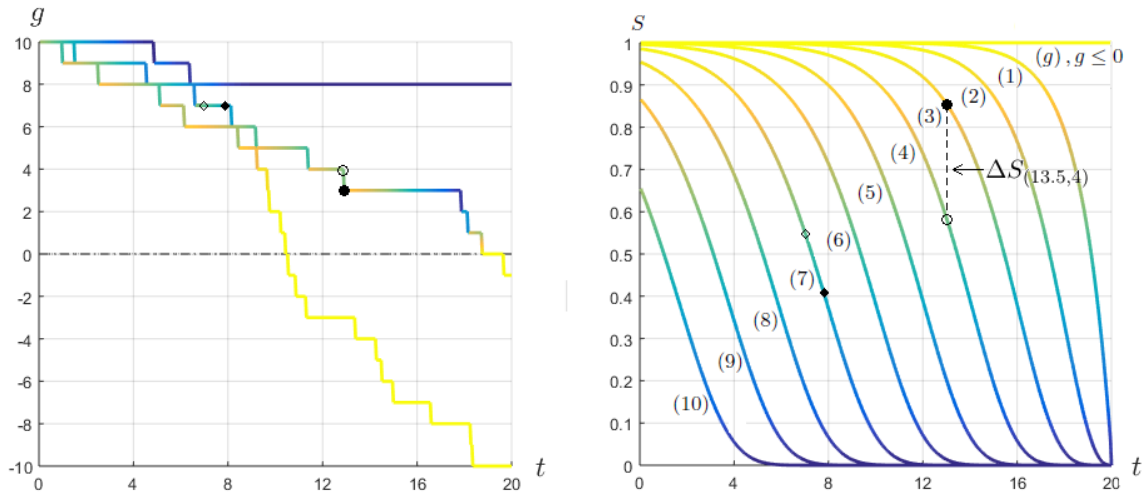
(a) Three gap paths; colour indicates $S_{(t,g_t)}{}^*$ (b) $S_{(t,g)}$ with fixed gap indicated by (g)

Figure 2.1: Time profiles of g and S; $F(c) = \frac{c}{q}, c \in [0, q]$, $g_0=10$, $(\tau,\lambda,q) = (20,0.95,0.8)$.*

*By Markov property, path likelihoods after crossing points are independent of prior paths. $S_0=0.65$.

Panel (a) depicts three simulated paths of $S_{(t,g_t)}$. The highest path shows a failing project and the darkening blue colour reflects the increasingly low success prospects with $S_{(t,g_t)}$ nearing zero by $t = 10$: almost all positive cost types choose **A** since gap $g_{10} = 8$ with $\tau - 10 = 10$ units of time left is near hopeless. At the other extreme, the path that ends up lowest becomes

increasingly yellow, reflecting nearly and then fully, guaranteed success as g nears 0 with plenty of time left and then falls below 0.

Panel (b) presents curves of $S_{(t,g)}$ as t varies for each fixed gap g. The empty and solid diamonds in (a) at $g = 7$ correspond to those in (b) where $S_{(t,g_t)}$ slides down $S_{(t,7)}$, marked by (7). By contrast, the later move from empty to solid circle at $t = 13.5$ in (a) depicts a drop in g from 4 to 3, causing S to jump up by pivotality $\Delta S_{(13.5,4)}$ equal to the dashed vertical distance from curve (4) to (3) in panel (b). Initial pivotality is about $0.85 - 0.65 = 0.2$ from curves (10) and (9) at $t = 0$ in (b) since $g_0 = 10$. At any *fixed* gap $g \geq 2$, $\Delta S_{(t,g)}$ first increases and then decreases over time. It usually falls to 0 as most campaigns are either clearly failing or clearly succeeding towards the end but $\Delta S_{(t,g_t)}$ can stay high if g_t falls at a specific intermediate rate. Indeed, pivotality takes its largest possible value $\Delta S_{(\tau,g_\tau)} = 1$ if $g_\tau = 1$ as shown by the vertical difference between curves (0) and (1) at $t = \tau = 20$. This does not contradict decreasing average pivotality because the probability of this cliff-hanger case is very low and at other gaps pivotality falls to zero.

2.3.4 The decreasing pivotality effect on expected bidding

From Eq. (2.6) and the fact that F is an increasing function, decreasing pivotality implies a tendency for bidding to fall. $S_{(t,g_t)}$ is a martingale so $S^{\text{bid}}_{(t,g_t)}$ falls at the same rate as pivotality $\Delta S_{(t,g_t)}$ in expectation. We call the effect of the expected fall in $S^{\text{bid}}_{(t,g_t)}$, equivalent to that in $\Delta S_{(t,g_t)}$, from t to x on the inspection rate of an imminent arrival, the (expected) *decreasing pivotality effect* (DPE), defined by

$$\mathcal{E}_x^{(t,g)} \triangleq F\left(\mathbb{E}_{(t,g)}\left(qS^{\text{bid}}_{(x,g_x)}\right)\right) - F\left(qS^{\text{bid}}_{(t,g)}\right) \tag{2.14}$$

The effect on the bidding rate is $\mathcal{E}_x^{(t,g)}$ times λq. For any gap g, $\mathcal{E}_x^{(t,g)} \leq 0$ if $t \leq x$ because $\mathbb{E}_{(t,g)}\left(qS^{\text{bid}}_{(x,g_x)}\right) \equiv q\left(S^{\text{bid}}_{(t,g)} + \mathcal{D}_x^{(t,g)}\right)$, $\mathcal{D}_x^{(t,g)} \leq 0$ from Eq. (2.13) and F is increasing. If F is differentiable, since $\mathcal{D}_{t+\text{d}t}^{(t,g)} = \mathcal{D}^{(t,g)}\,\text{d}t + O(\text{d}t^2)$, the rate of change DPE is

$$\mathcal{E}^{(t,g)} \triangleq \lim_{\text{d}t \downarrow 0} \left(\frac{1}{\text{d}t}\left(\mathcal{E}_{t+\text{d}t}^{(t,g)}\right)\right) = qF_c(qS^{\text{bid}}_{(t,g)})\mathcal{D}^{(t,g)} \tag{2.15}$$

On top of this gradual reduction, any discontinuities in F cause bidding $\beta_{(t,g_t)}$ to drop by discrete amounts at a set of critical dates. The fall at a critical date is λq times the probability mass of the types just willing to inspect at that date as we detail below.

Discontinuities. Let CDF F have a mass z^k atom at $c = c^k$ for each $k \in \{0, 1, 2, ..., K\}$, indexed so that $0 = c^0 < c^1 < c^2 < ... < c^K$; z^0 can be 0 and K could be infinite. Then for any g and $c^k < q$, by Lemma 2.1 and

the intuitive fact that $S_{(t,g)}^{\mathrm{bid}}(\tau) \to 1$ as $\tau \to \infty$, a unique duration $\hat{\tau}_g^k$ satisfies $S_{(0,g)}^{\mathrm{bid}}(\tau) = c^k/q$. For $c^K = q$, $g \geq 2$, we write $\hat{\tau}_g^K = \infty$. $\hat{\tau}_g^k = 0$ if $g \leq 1$ but a strictly positive duration is needed for higher gaps if $k > 0$. At $t = \hat{t}_g^k \triangleq \tau - \hat{\tau}_g^k$, $\mathcal{E}_{t_+}^{(t,g)} \triangleq \lim_{dt\downarrow0}\left(\mathcal{E}_{t+dt}^{(t,g)}\right) = -z^k$. Sections 2.4.5 and 2.4.6 illustrate.

2.3.5 Expected bid dynamics and the Jensen effect

We apply (GEN) to the stochastic process $\beta_{(t,g_t)}$ to reveal expected changes in the bid rate from a generic state (t,g). Where F is continuous, $\beta_{(t,g_t)}$, like $S_{(t,g_t)}^{\mathrm{bid}}$, falls gradually as time passes without bids (the direct negative effect of time, $\dot{\beta}_{(t,g)} < 0$) and features positive jumps $\Delta\beta_{(t,g)} \equiv \beta_{(t,g-1)} - \beta_{(t,g)}$ whenever bids occur (the indirect positive effect). Since $\beta_{(t,g)}$ is proportional to $F\left(qS_{(t,g)}^{\mathrm{bid}}\right)$ with positive multiplier (λq), the dynamics are qualitatively similar but there are two differences. First, discontinuities in the cost distribution F introduce discrete bid rate drops at the dates $\hat{t}_g^k$ noted above; smoothly decreasing pivotality (DP) has discrete effects (DPE). Second, in contrast to the unambiguous decreasing average profile for $\Delta S_{(t,g_t)}$ and $S_{(t,g)}^{\mathrm{bid}}$ (DP), the direct negative effect of time does not always dominate its indirect positive effect via bids. The average effect of time now depends on the degree of uncertainty and the shape of the cost distribution.

Over an infinitesimal time interval, either one bid arrives or none do (Fig. 2.2 adds details below). A bid arrival is good news, raising S^{bid}, while no bid is bad news, lowering S^{bid}. We say that the good news outweighs the bad news if the probability that an arrival inspects at $(t + \epsilon, g_{t+\epsilon})$ exceeds the probability at (t,g). This occurs if the density of inspection costs just above $qS_{(t,g)}^{\mathrm{bid}}$ is greater than the density just below it. Increasing density is equivalent to convexity of the CDF. In economic terms, there are increasing returns to good news: bidders with higher inspection costs who become willing to inspect after good news outweigh those with lower inspection costs who cease to inspect after bad news. In the opposite case where F is concave, the returns are decreasing. Formally, we define the Jensen effect (JE) as the impact of uncertainty, about bidding on (t,x), on the inspection probability of a bidder at $x > t$ as anticipated from state (t,g),

$$\mathcal{J}_x^{(t,g)} \triangleq \mathbb{E}_{(t,g)}\left(F\left(qS_{(x,g_x)}^{\mathrm{bid}}\right)\right) - F\left(\mathbb{E}_{(t,g)}\left(qS_{(x,g_x)}^{\mathrm{bid}}\right)\right) \tag{2.16}$$

By Jensen's inequality, this is positive if F is convex, negative if concave and zero if affine. We also define a rate of change JE as the expected rate of change in inspections,

$$\mathcal{J}_{(t,g)} \triangleq \lim_{dt\downarrow0}\left(\frac{1}{dt}\left(\mathcal{J}_{t+dt}^{(t,g)}\right)\right) = \beta_{(t,g)}\left(\Delta F(qS_{(t,g)}^{\mathrm{bid}}) - F_c(qS_{(t,g)}^{\mathrm{bid}})(q\Delta S_{(t,g)}^{\mathrm{bid}})\right) \tag{2.17}$$

The second equality, derived in Appendix 2.A proof of Proposition 2.4, assumes F is differentiable. It is useful to predict magnitudes. It reconfirms that JE is positive if F is

convex since a convex graph lies above its tangent while JE is negative if F is concave.

Time's overall average impact sums these effects: from Eq. (2.6), rescaling by λq gives the change in bidding expected by date $x \geq t$ from state (t, g),

$$\mathbb{E}_{(t,g)}\left(\beta_{(x,g_x)}\right) - \beta_{(t,g)} \equiv \lambda q \left(\mathcal{J}_x^{(t,g)} + \mathcal{E}_x^{(t,g)}\right) \tag{2.18}$$

If F is continuous, dividing by ϵ and letting $\epsilon \to 0$ gives the decomposition,

$$\mathcal{L}_{(t,g)}^{\beta} \equiv \lambda q \left(\mathcal{J}_{(t,g)} + \mathcal{E}_{(t,g)}\right) \tag{2.19}$$

Since DPE $\mathcal{E}_x^{(t,g)}$ is always weakly negative, bidding is expected to fall if JE $\mathcal{J}_x^{(t,g)}$ is either small or neutral, as when F is affine, or negative as when F is concave. In the last case, the Jensen and decreasing pivotality effects push downwards together: $\mathcal{J}_x^{(t,g)} \leq 0$ and $\mathcal{E}_x^{(t,g)} \leq 0$, so by Eq. (2.18),

Proposition 2.3. *A flat or decreasing inspection cost density on $[0, q]$ (weakly concave F) generates expected bid rates that fall from any state (t, g): $\forall x > t, \mathbb{E}_{(t,g)}\left(\beta_{(x,g_x)}\right) \leq \beta_{(t,g)}$.*

The uniform distribution is a special case (see Section 2.4.1). When $F_c \equiv 1/q$ on cost range $[0, q]$, F is linear, β and S^{bid} are exactly proportional and $\mathbb{E}_{(t,g)}\left(\beta_{(x,g_x)}\right) - \beta_{(t,g)} = \lambda q \, \mathcal{E}_x^{(t,g)} \leq 0$. JE is null and $\mathcal{L}_{(t,g)}^{\beta} = \lambda q \mathcal{E}_{(t,g)} = \lambda q \mathcal{D}_{(t,g)} \leq 0$.

When instead F is convex, the increasing returns to good news, positive JE, counteracts the decreasing pivotality effect (DPE) and potentially causes the expected bid rate to rise. Before seeking a precise condition for this, Fig. 2.2 illustrates how it can happen.

Over infinitesimal time period ϵ, a bid arrives with probability $\epsilon \beta_{(t,g)}$, making $\tilde{g}_{t+\epsilon} = g - 1$ so $S_{(t,g_t)}^{\mathrm{bid}}$ jumps up to $S_{(t+\epsilon,g-1)}^{\mathrm{bid}}$ and otherwise no bid arrives (probability $1 - \epsilon \beta_{(t,g)}$) and the bid-contingent success rate $S_{(t,g)}^{\mathrm{bid}}$ moves to $S_{(t+\epsilon,g)}^{\mathrm{bid}}$. The DP arrow indicates the expected fall in this success prospect: $\mathbb{E}_{(t,g)} S_{(t+\epsilon,\tilde{g}_{t+\epsilon})}^{\mathrm{bid}} \leq S_{(t,g)}^{\mathrm{bid}}$ as $\mathcal{D}_{t+\epsilon}^{(t,g)} \leq 0$. The DPE arrow shows how this reduces the expected incentive to inspect at $t + \epsilon$ by $-q \mathcal{D}_{t+\epsilon}^{(t,g)}$, reducing expected inspections by $-\mathcal{E}_{t+\epsilon}^{(t,g)} \approx -q F_c \left(S_{(t,g)}^{\mathrm{bid}}\right) \mathcal{D}_{t+\epsilon}^{(t,g)}$. Finally, the JE arrow indicates the impact of the variance in $S_{(x,g_x)}^{\mathrm{bid}}$ at $x = t + \epsilon$, $\mathcal{J}_{t+\epsilon}^{(t,g)} = \mathbb{E}_{(t,g)} F(q S_{(t+\epsilon,\tilde{g}_{t+\epsilon})}^{\mathrm{bid}}) - F\left(\mathbb{E}_{(t,g)}\left(q S_{(t+\epsilon,g_{t+\epsilon})}^{\mathrm{bid}}\right)\right)$. In the figure, F is sufficiently convex to create a positive JE that exceeds the DPE. The generator $\mathcal{L}_{(t,g)}^{\beta}$ of $\beta_{(t,g_t)}$ provides a precise *sufficient* condition.

For affine and quadratic F, the proof of Proposition 2.4 reveals that

$$\mathcal{L}_{(t,g)}^{\beta} = \lambda q \left[q F_c \left(q S_{(t,g)}^{\mathrm{bid}}\right) \mathcal{L}_{(t,g)}^{S^{\mathrm{bid}}} + \frac{q^2}{2} F_{cc}\left(q S_{(t,g)}^{\mathrm{bid}}\right) \nu_{(t,g)}^{S^{\mathrm{bid}}}\right] \tag{2.20}$$

where $\nu_{(t,g)}^{S^{\mathrm{bid}}} \triangleq \lim_{dt\downarrow 0}\left(\frac{1}{dt}\left(\mathbb{V}_{(t,g)}\left(dS_{(t,g_t)}^{\mathrm{bid}}\right)\right)\right) = \beta_{(t,g)}\left(\Delta S_{(t,g)}^{\mathrm{bid}}\right)^2$ is $S_{(t,g_t)}^{\mathrm{bid}}$'s jump variance given $g_t = g$. As $\mathcal{L}_{(t,g)}^{S^{\mathrm{bid}}} = \mathcal{L}_{(t,g)}^{\Delta S}$, the first term in Eq. (2.20) represents the expected impact of decreasing pivotality. The second term is the Jensen effect, driven by variance and convexity: $\mathcal{J}_{(t,g)} = \frac{1}{2}\,\beta_{(t,g)}\left(q\Delta S_{(t,g)}^{\mathrm{bid}}\right)^2 F_{cc}$.
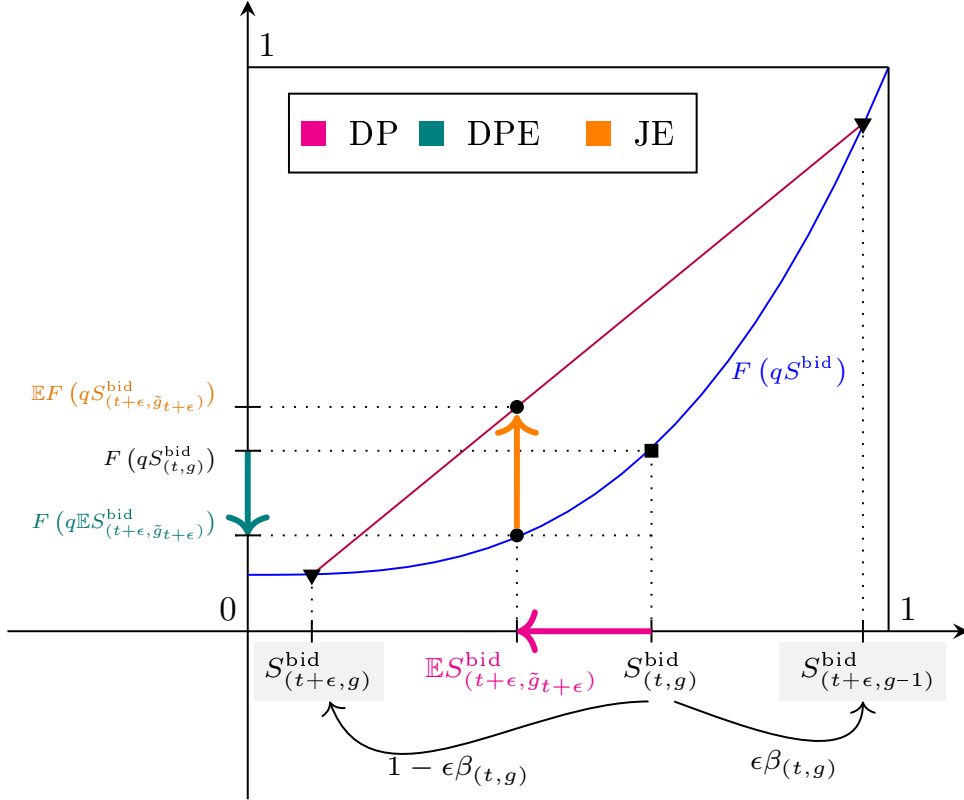
Figure 2.2: Decreasing pivotality DP, effect DPE and Jensen effect JE for convex CDF.*

*Curved arrows indicate state-transition probabilities.
All expectations condition on (t,g): $\mathbb{E}$ is shorthand for $\mathbb{E}_{(t,g)}$.

Beyond the quadratic case, we prove that $\widetilde{\mathcal{L}}^{\beta}_{(t,g)} \triangleq \mathrm{RHS}(2.20)$ is a lower bound on the bid rate generator for any power CDF with exponent two or higher to conclude:

Proposition 2.4. *Sufficient uncertainty and convexity guarantee a rising expected bid rate: for a polynomial F with non-negative coefficients on powers two and above, $\mathcal{L}^{\beta}_{(t,g)} > 0$ if*

$$\frac{q}{2} F_{cc}\left(q S^{\mathrm{bid}}_{(t,g)}\right) \nu^{S^{\mathrm{bid}}}_{(t,g)} > F_c\left(q S^{\mathrm{bid}}_{(t,g)}\right) \left| \mathcal{L}^{S^{\mathrm{bid}}}_{(t,g)} \right| \tag{2.21}$$

Proof *in Appendix 2.A.* Again, uncertainty and convexity are both needed for an increasing slope. We conclude the analytic illustration by substituting for $\nu^{S^{\mathrm{bid}}}_{(t,g)}$ when F is a power ρ function. Then $\frac{cF_{cc}}{F_c} = \rho - 1$ and inequality (2.21) is also necessary so that the necessary and sufficienty condition for a positive slope at (t,g) is

$$\frac{1}{2} \frac{\Delta S^{\mathrm{bid}}_{(t,g)}}{S^{\mathrm{bid}}_{(t,g)}} (\rho - 1) > \frac{\Delta \beta_{(t,g)}}{\beta_{(t,g)}} \tag{2.22}$$

2.3.6 Average bidding and the state transition process

To study average bidding at time t after initial state $(0, g_0)$, we characterize the probability weights over possible g_t at t. The *average bid rate,*

$$A_t = \sum_{g=-\infty}^{g_0} Q_{(t,g)} \beta_{(t,g)} \tag{2.23}$$

where $Q_{(t,g)} \triangleq Q_{(t,g)}^{(0,g_0)}$ and $Q_{(t,g)}^{(t',g')}$ denotes the transition probability from (t',g') to (t,g). The bidding profile is the plot of A_t against time. We aim to characterize its slope. $\dot{A}_t \triangleq \frac{\partial}{\partial t}(A_t) \equiv \frac{\mathrm{d}}{\mathrm{d}t}(A_t)$ since A only depends on t. To derive $Q_{(t,g)}$, we solve recursively for the full set of transition matrices $Q_{(t,g)}^{(t',g')}$.

Lemma 2.2. *For any $t \geq t'$, the transition process is characterized by*

$$Q_{(t,g)}^{(t',g')} = 0 \text{ for all } g > g'; \qquad Q_{(t,g)}^{(t',g)} = \exp\left(-\int_{t'}^{t} \beta_{(x,g)} \,\mathrm{d}x\right) \tag{2.24}$$

$$Q_{(t,g)}^{(t',g')} = \int_{t'}^{t} \exp\left(-\int_{t'}^{T} \beta_{(x,g')} \,\mathrm{d}x\right) \beta_{(T,g')} Q_{(t,g)}^{(T,g'-1)} \,\mathrm{d}T, \text{ for } g \leq g'-1 \tag{REC-Q}$$

Proof.   Since the gap cannot increase, for any $t \geq t'$, $Q_{(t,g)}^{(t',g')} = 0$ for all $g > g'$. For $g = g'$, the transition probability equals the probability of no bid on (t', t): $Q_{(t,g)}^{(t',g)} \equiv n_t^{(t',g)}$ of Eq. (2.7) gives the second equality in Eq. (2.24). The recursive step parallels Eq. (REC-S): for any $g \leq g'-1$, conditioning on (t', g') and on the first stopping time T after (t', g') with density $n_T^{(t',g')} \beta_{(T,g')}$, the Law of Iterated Expectations yields Eq. (REC-Q). ∎ Appendix 2.C.3 adds an alternative derivation of (REC-Q) via the ODE for $Q_{(t,g)}^{(t',g')}$ that varies t'. For use in the final subsection, it also derives the adjoint ODE by varying t:

$$\dot{Q}_{(t,g)} = Q_{(t,g+1)} \beta_{(t,g+1)} - Q_{(t,g)} \beta_{(t,g)}, \tag{ODE-Q}$$

Intuitively, the rate of change in the probability of gap $g_t = g$ equals the probability of reaching this gap via a bid from state $(t_-, g+1)$ minus the probability that the gap falls below g via a bid in state (t_-, g).

Conditioning

For profiles conditioned on campaign success or failure, we restrict the space of paths to those ending with $g_{\tau_+} \leq 0$ or $g_{\tau_+} > 0$, respectively. A path passing through state (t, g) with a bid at that moment ends in success with probability $S_{(t,g-1)}$. So conditioning bid intensity at t on success, rescales the averaging probabilities by $S_{(t,g-1)}/S_0$. Similarly, conditioning on failure rescales by $(1 - S_{(t,g-1)})/(1 - S_0)$.

$$A_t^{\mathcal{S}} = \frac{1}{S_0} \sum_{g=-\infty}^{g_0} Q_{(t,g)} \beta_{(t,g)} S_{(t,g-1)} \tag{25-$\mathcal{S}$}$$

$$A_t^{\mathcal{F}} = \frac{1}{1 - S_0} \sum_{g=-\infty}^{g_0} Q_{(t,g)} \beta_{(t,g)} \left(1 - S_{(t,g-1)}\right) \tag{25-$\mathcal{F}$}$$

Conditioning on success shifts the distribution of gaps g at t towards lower gaps. This implies a weighting bias towards higher bid rates $\beta_{(t,g)}$. There is no scope for such gap selection at $t = 0$ where $Q_{(t,g)} = 0$ for all $g \neq g_0$, so this selective weighting effect initially increases with time. Success conditioning also positively selects for bidding at t. This bid-selection effect may also increase over time as on campaigns that are otherwise close to failing. At $t = 0$ where only the bid-selection bias is operative, success-conditioning multiplies the bid rate by one plus relative pivotality:

$$A_0^{\mathcal{S}}/A_0 = \frac{\beta_{(0,g_0)} S_{(0,g_0)}^{\text{bid}}}{S_0} \Big/ \beta_{(0,g_0)} = 1 + \frac{\Delta S_{(0,g_0)}}{S_{(0,g_0)}}$$

On average, absolute pivotality falls over time but relative pivotality is high in low success rate states creating an opposite trend. The general expressions are complicated but for intuition notice how time increasingly polarizes the bid-contingent success rate into values close to zero for relatively high gaps and close to one for relatively low gaps. That rate is the weighting factor when conditioning on success. Its negative monotonicity in g combined with the same monotonicity of the bid rate can generate increasing differences between the success and failure conditioned bid profiles.

Since conditioning on failure features the opposite selection biases, these effects can explain why the conditional average bid profiles fan out relative to the unconditional profile in almost all our figures for the examples of the next section. However, the general comparisons are more complicated and Fig. 2.D.2a provides an exception to the pattern.

2.3.7 The slope of the average bid profile

The slope of the bid profile can be expressed as a weighted average of $\mathcal{L}_{(t,g)}^{\beta}$.

Lemma 2.3. *The time gradients of (i) unconditional, (ii) $\mathcal{S}$-, $\mathcal{F}$-conditional average bid rates are*

$$\dot{A}_t = \sum_{g=-\infty}^{g_0} Q_{(t,g)} \mathcal{L}_{(t,g)}^{\beta}$$

$$\dot{A}_t^{\mathcal{S}} = \frac{1}{S_0} \sum_{g=-\infty}^{g_0} Q_{(t,g)} \mathcal{L}_{(t,g)}^{\beta} S_{(t,g-1)}, \quad \dot{A}_t^{\mathcal{F}} = \frac{1}{1-S_0} \sum_{g=-\infty}^{g_0} Q_{(t,g)} \mathcal{L}_{(t,g)}^{\beta} \left(1 - S_{(t,g-1)}\right) \tag{2.26}$$

Proof of Lemma 2.3. (i) Differentiating A_t and using (ODE-Q) of the Markov process,

$$
\begin{aligned}
\dot{A}_t &= \sum_{g=-\infty}^{g_0} \left(\dot{Q}_{(t,g)} \beta_{(t,g)} + Q_{(t,g)} \dot{\beta}_{(t,g)} \right) \\
&= \sum_{g=-\infty}^{g_0} \left(\left(Q_{(t,g+1)} \beta_{(t,g+1)} - Q_{(t,g)} \beta_{(t,g)} \right) \beta_{(t,g)} + Q_{(t,g)} \dot{\beta}_{(t,g)} \right) \\
&= \sum_{g=-\infty}^{g_0} Q_{(t,g)} \left(\beta_{(t,g)} \left(\beta_{(t,g-1)} - \beta_{(t,g)} \right) + \dot{\beta}_{(t,g)} \right) = \sum_{g=-\infty}^{g_0} Q_{(t,g)} \mathcal{L}_{(t,g)}^{\beta} \tag{2.27}
\end{aligned}
$$

since $\sum_{g=-\infty}^{g_0} Q_{(t,g+1)}\beta_{(t,g+1)}\beta_{(t,g)} = \sum_{g=-\infty}^{g_0} Q_{(t,g)}\beta_{(t,g)}\beta_{(t,g-1)}$ given that $Q_{(t,g_0+1)} = 0$.

(ii) We derive the conditional results similarly in Appendix 2.A. ∎

By this lemma, an everywhere positive $\mathcal{L}^{\beta}_{(t,g)}$ is sufficient for a monotone increasing bid profile, and everywhere negative generators guarantee monotone decreasing profiles.

Corollary 2.2. *If $\mathcal{L}^{\beta}_{(t,g)} \geq 0$ for all (t,g), the aggregate and conditional bid profiles are increasing over time. In the opposite case of $\mathcal{L}^{\beta}_{(t,g)} \leq 0, \forall (t,g)$, they are decreasing.*

Remark 2.1. *$\mathcal{L}^{\mathcal{L}^{\beta}}_{(t,g)} \geq 0, \forall (t,g)$ implies a convex bid profile; concave if $\mathcal{L}^{\mathcal{L}^{\beta}}_{(t,g)} \leq 0, \forall (t,g)$.*

Applying Corollary 2.2 to Propositions 2.3 and 2.4, delivers two results:

Proposition 2.5. *A weakly concave CDF $F(c)$ generates a weakly decreasing average bid profile: $\dot{A}_t \leq 0, \forall t$. Strict concavity implies a strictly negative slope if $g_0 \geq 2$.*

The strict claim uses two facts: $Q_{(t,g_0)} > 0$ for any t and if the cost distribution has full support, Lemma 2.1 holds with strict inequalities for any $g_0 \geq 2$.

Proposition 2.6. *Imposing Proposition 2.4's convexity and uncertainty conditions (2.21) at all $g \geq 2$ guarantees a strictly increasing average bid profile, $\dot{A}_t > 0, \forall t$, if $g_0 \geq 2$.*

A weaker condition for a strict positive slope is possible: relax the conditions of Proposition 2.4 to allow equality except for holding strictly at all $(t, g(t))$ where $g(t)$ is any function of t with $Q_{(t,g(t))} > 0$ for all t. See also Lemma 2.6.

2.4   Canonical distribution classes

In this section, we apply our results to specific functional forms of the CDF $F(c)$ of inspection costs. We begin with linear, quadratic and higher power distributions. Linear and more generally, affine CDFs, correspond to uniform distributions and preclude any Jensen effect (JE). As a result, the decreasing pivotality effect (DPE) perfectly explains the shape of the negative sloped bidding profile. Quadratic and generic power CDFs show how positive and negative Jensen effects influence bidding profiles. Generic single-peaked distributions guarantee a unique signed Jensen effect if the mode lies outside the relevant cost range, as we illustrate. Distributions with atoms are neither concave nor convex. Atoms cause bid rate discontinuities but only at a set of critical dates with zero measure. The central message is that our generator-based analysis continues to apply at all other dates. The specific result is that an atom at the limit cost q combined with any moderate continuous distribution readily generates a U-shape.

While somewhat special, the homogenous case is a good starting point because we can develop a simplified characterization. Average pivotality falls smoothly but its effect on

bidding occurs exclusively at critical dates and in the form of discrete drops. Specific to pure homogeneity, the JE is zero except at critical dates and the DPE dominates there. More generally, the JE is positive for gaps just above a critical gap thanks to bidders with inspection costs below the atom. This is because the good news of a bid from a lower cost bidder heats up the campaign by causing the atom of bidders to start inspecting. Generic binary cost distributions generate saw-tooth bid profiles: discrete drops from the DPE intersperse smooth upward slopes from the positive JE. If the high cost type is near to the limit cost q, this positive JE is negligible until late in the campaign. In contrast, decreasing pivotality is stronger early in the campaign. These factors rise to a U-shape. The final subsection builds on this insight by demonstrating a fully smooth U-shape from a uniform distribution with an atom at q. Any decreasing, flat or moderate increasing density would serve equally well for the U-shape. The atom at or near q is important and we motivate a limit cost atom below.

The central logic behind the U-shape derives from information revelation. As time passes, it usually becomes clear which projects will succeed and which will fail. Pivotality is then usually small at both extremes of this bifurcation. On campaigns that are pretty clearly succeeding, success prospects are already close to unity and cannot go higher. On campaigns that seem headed for failure, one bid near the deadline has little chance of making the difference. So pivotalities tend to bunch up near zero in the late phase of a campaign. This confines any substantial negative slopes from the DPE to earlier on. The fact that success prospects tend to bifurcate into rates close to one and rates close to zero as the deadline approaches has a second consequence: when rates are close to one, this activates bidder types with costs near the upper limit q. A high density of such types then generates an upward bid slope in the late phase.

As justified in Section 2.2, we again take $F(c)$ restricted to support $[0, q]$. Appendix 2.B.1 formally shows how truncating a generic CDF by removing values above q and reducing the arrival rate to $\lambda' = \lambda F(q)$, then substituting values below 0 with an atom of size $F(0)$ at 0 generates identical outcomes to the original CDF with arrival rate λ. This motivates an atom of bidders at zero. We provide comparative static effects on slope magnitudes of shifting this zero atom, denoting its mass by z^0. The zero atom represents the potentially identifiable set of bidders who already know their taste for the entrepreneur's product *or* have negative inspection costs, perhaps because they are fans, contacts or friends.

The truncation argument does not justify an atom at q since bidders strictly above q simply reduce the effective arrival rate. Nonetheless, an atom at q is equivalent to a plausible bidder type that we describe now. In addition to the bidders described in the model, we consider a set of bidders who follow a rule-of-thumb. They save on thinking costs by only thinking about available prospects, rather than also assessing hypothetical ones. For gaps above 2, the option to buy is contingent on an uncertain funding outcome

because even if the bidder who arrives chooses to bid, he relies on enough others later choosing to bid as well. The rule of thumb bidders turn away from such uncertain prospects and dedicate their attention to other options such as direct offers of goods. Since bidders with $c = q$ are only willing to inspect, choose substrategy **C**, when the gap $g_t \leq 1$, it is immediate that rule-of-thumb bidders follow the exact same strategy so this justifies assuming an atom of bidders at q with size equal to the fraction of bidders who follow the rule.

Notice that this rule of thumb does not discriminate between degrees of uncertainty. That makes sense because it may require thought to compute that degree. Indeed, the analysis for computing the conditional success rates in this paper is computationally quite complicated. While bidders can be assumed to form approximate expectations based on experience, even that can be quite difficult. In addition, some simple approximations, such as viewing the bidding option as an approximately sure trade option when the gap is below say 5% of the initial threshold, generate very similar profiles to the proposed heuristic. One other approximation bears mention. A reasonable variant of the heuristic bidder only contemplates products that have already succeeded. This approximation only distinguishes projects that will definitely result in production, *independent* of whether the individual bidder who is evaluating, ends up choosing to bid, from projects where there is uncertainty. Such bidders would require $g \leq 0$ but the profile implications are extremely similar, just with the upward slope pushed further towards the end of the campaign duration.

Another motivation for this atom at q comes from a distinct form of procedural rationality. Suppose bidders are impulsive once they begin to evaluate a prospect. They know that if they consider an innovative product and discover that they like it, then they will really want to have it. This impulse to buy implies a risk of feeling dissatisfied. So, considering the product before its success is guaranteed, risks creating a strong desire that provokes a pain of disappointment or dissatisfaction in the event that the project fails to reach its threshold. In the model, we assumed that the bidder loses his inspection cost c and faces no such pain in the event that a project fails to get funded. If instead failure imposes pain on those who inspect and like the good, inspecting can be very attractive for a sure offer, as when $g \leq 1$ and c_t is small or negative, even though inspecting is a bad idea if the project may fail. We denote this pain term by c_R since the bidder regrets having inspected if the campaign fails. It differs from the baseline inspection cost c because it is not forfeited in the event of a project success. If c_R is large enough, it is easily seen that group 2 bidders only inspect when $g \leq 1$ so they are again equivalent to baseline bidders with $c = q$.

The analysis of comparative statics in the canonical cases that follow use a couple of intuitive results about the effects of parameters and first-order-stochastic-domination which we denote by FOSD: $F(\cdot) \underset{\text{FOSD}}{\succeq} F'(\cdot)$ if $F'(c) \geq F(c)$ for all $c \in [0, q]$.

Lemma 2.4. *If* $F(\cdot) \underset{FOSD}{\succeq} F'(\cdot)$ *then* $S_{(t,g)}(F') \geq S_{(t,g)}(F)$ *for all* t, g.

Intuitively, high costs dissuade inspection so they lower success rates. The inductive **proof** (in Appendix 2.A) shows that *FOSD* raises S by raising the probability of a bid given S^{bid} and S^{bid} is taken to satisfy the hypothesis. We later use the corollary that a proportionate shift in probability from $c > 0$ values onto the zero-type uniformly raises the success rate. Finally, Eq. (2.9) within the inductive proof of Lemma 2.4 proves

Lemma 2.5. λ *and* q *both increase* $S_{(t,g)}$ *for all* t, g.

2.4.1 Affine CDF

If inspection costs follow a uniform distribution with atom z at $c = 0$, the CDF is affine (linear if $z = 0$):

$$F(c) = z + (1 - z)\left(\frac{c}{q}\right) \tag{2.28}$$

$F_{cc} \equiv 0$, so the Jensen effect is null. With decreasing pivotality as the only force, the bid profile has a negative slope. The bid rate generator for the quadratic CDF applies to the affine CDF, so setting $F_c = \frac{1-z}{q}$ in (2.20) and recalling $\mathcal{L}^{S^{\mathrm{bid}}}_{(t,g)} = \mathcal{L}^{\Delta S}_{(t,g)}$,

$$\mathcal{L}^{\beta}_{(t,g)} = \lambda q (1 - z) \mathcal{L}^{\Delta S}_{(t,g)} \quad \leq 0 \tag{2.29}$$

Fig. 2.3 exhibits the downward bid slopes for averages across all projects in black, conditioned on success (in green) and failure (in red). Notice how the unconditional bid profiles (in black) in panels (a) and (b) are identical, subject to rescaling, to the corresponding average pivotality profiles of (c) and (d). So the decreasing bid profiles reflect the Decreasing Pivotality Effect (DPE) operating in isolation.

Panels (e) and (f) add pivotality profiles for $g_0 = 16$ and 22 so that g_0 rises from 16 to $18, 20, 22$ moving clockwise from panel (e) to (c), (d), (f). This illustrates how average pivotality and the equivalent bidding profiles become increasingly convex as g_0 rises. For $g_0 = 16$ and 18, the pivotality profile is initially concave and becomes convex as the deadline approaches. For low initial gaps, the fixed-g success curves of Fig. 2.1b are bunched up near the unit upper bound and concave so that pivotality, equal to the vertical difference, is moderate and initially falls slowly but accelerates with time, generating initial concave regions for average pivotality. In our example, $g_0 = 16$ of (e) is a low initial gap; $S_0 = S_{(0,16)}$ is high at 0.94 whereas $S_{(0,18)} = 0.76$. The curves in Fig. 2.1b inflect from concave to convex. Similarly, average pivotality shrinks and becomes convex as pivotality becomes increasingly constrained by the zero lower bound; pivotality ends up at zero in the last moments except in knife-edge cases with a gap of exactly 1. When the initial gap is high, as with $g_0 = 22$ of panel (f), the pivotality curve starts in its convex phase because time is already too scarce; $S_{(0,22)} \approx 0$. Though pivotalities are scaled down towards zero, (f) has the highest curvature of the four.
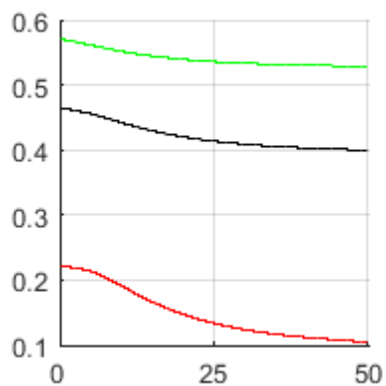
(a) Bid profile for $g_0 = 18$
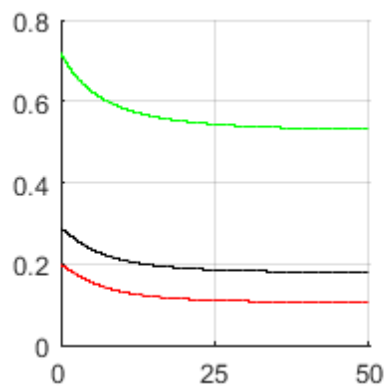
(b) Bid profile for $g_0 = 20$
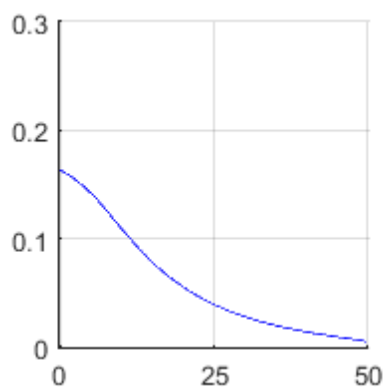
(c) pivotality for $g_0 = 18$
(S_0=0.70)
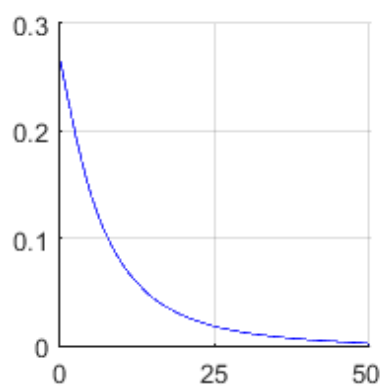
(d) Pivotality for $g_0 = 20$
(S_0=0.17)
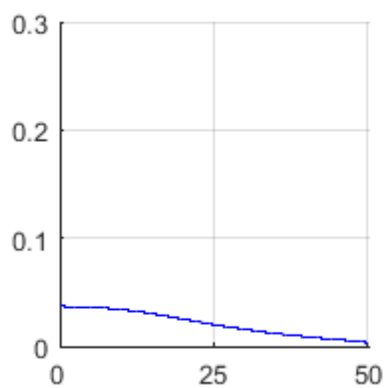
(e) Pivotality for $g_0 = 16$
(S_0=0.94)
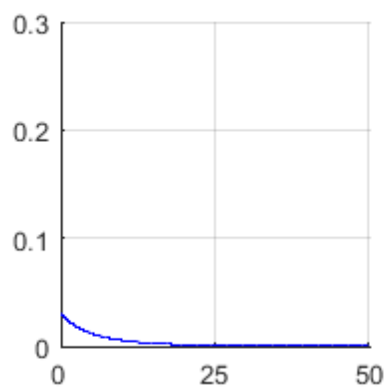
(f) Pivotality for $g_0 = 22$
(S_0=0.0051)

Figure 2.3: Profiles of bids —$A_t^{\mathcal{S}}$ —A_t —$A_t^{\mathcal{F}}$ and pivotality —$A_t^{\Delta S}$ against time t for a linear CDF with $z = 0.2$ given $g_0 \in \{16, 18, 20, 22\}$ and $(\tau, \lambda, q) = (50, 0.7, 0.75)$

(a) $(z,q){=}(0,0.75)$ (b) $(z,q){=}(0.2,0.692)$ (c) $(z,q){=}(0.5,0.604)$
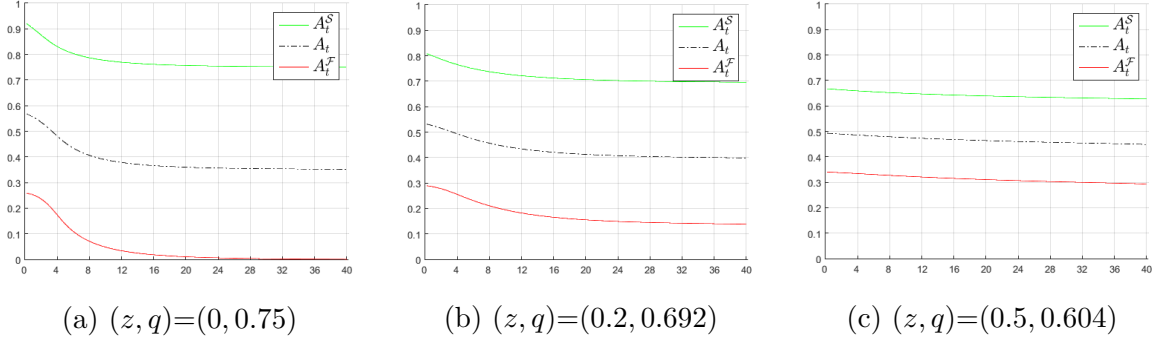
Figure 2.4: Comparative statics on z for linear CDF with q compensating to give $S_0 \approx 0.46$; $g_0 = 20$ and $(\tau, \lambda) = (40, 0.7, 0.75)$.

Conjecture 2.1. *The pivotality profile is less concave or more convex for higher $g_0/\lambda\tau$.*

Intuitively, early bids are critical on projects that set very high thresholds relative to time available and bidder arrival intensity. This clarifies the varying strength of decreasing pivotality.

As shown by Lemmas 2.4 and 2.5, both z and q raise S. Increasing z and q have the opposite qualitative effect to increasing $g_0/\lambda\tau$. The impact of z is not monotonic as shown by Fig. 2.D.2a: the downward average slope is most pronounced for the intermediate z values, specifically, $0.5, 0.6$. In the first four panels, $z \leq 0.4$ and the scarcity of the zero-cost, always-inspect bidders makes it hard for the campaign to get off the ground. Success is unlikely, reflected in the average profile A_t, in black, nearly coinciding with the failure-conditional curve $A_t^{\mathcal{F}}$, in red, with both quite flat and near zero. By contrast, the success-conditional bid profile $A_t^{\mathcal{S}}$ slopes downwards quite sharply because successes pretty much have to start well in these adverse cost settings. Returning to the average profile, as the zero cost atom rises from 0.6 to 1, the downward slope becomes increasingly flat as the strategically insensitive zero-cost types (who do not react to changes in S^{bid}) become increasingly prevalent. Dynamic effects vanish in the last panel with $z = 1$. In this linear setting, the effects are muted, but we find stronger parallel effects from higher power CDFs.

Appendix 2.B.3 provides explicit comparative statics of z on bid profiles in the case of $g_0 = 2$. The magnitude of the expression for the slope in Appendix 2.B.3 is increasing in t, so the profile is concave when $g_0 = 2$. Appendix 2.B.3 also proves that z has a consistent negative impact on slope magnitude but in general, z has two effects: raising success and reducing average bidder sensitivity to success. To isolate the strategic sensitivity effect, Fig. 2.4 illustrates the impact of changing z with opposite, compensatory changes in q to fix S_0. Raising z from panel (a) to (c) now always reduces the bidding slope's magnitude.

2.4.2 Quadratic CDF

A linear density with atom z at $c = 0$ produces the quadratic CDF,

$$F(c) = z + (1 - z)\left(\frac{c}{q}\right)^2 \tag{2.30}$$

This implies $F_c(qS_{(t,g)}^{\text{bid}}) = \frac{2(1-z)S_{(t,g)}^{\text{bid}}}{q}$ and $F_{cc}(qS_{(t,g)}^{\text{bid}}) = \frac{2(1-z)}{q^2}$. For $z < 1$ and $g > 1$, which ensure that $\nu_{(t,g)}^{S^{\text{bid}}} > 0$ and $F_{cc}(qS_{(t,g)}^{\text{bid}}) > 0$, the rising bid rate condition (2.21) simplifies to

$$z \geq \zeta_{(t,g)}(z) \triangleq 1 - \frac{1}{1 + S_{(t,g)}^{\text{bid}}(z)\left(2S_{(t,g-1)}^{\text{bid}}(z) + S_{(t,g)}^{\text{bid}}(z)\right)} \tag{2.31}$$

Sufficient mass z on zero-cost types generates enough bidding on those moribund campaigns in states with S close to 0 to create sufficient variance, via good news shocks, for a large Jensen effect (JE). Since $S^{\text{bid}} \leq 1$ always, RHS(2.31) is bounded above by 3/4 giving,

Lemma 2.6. *Quadratic CDF* (2.30) *generates a strictly increasing bid profile if* $z \in (3/4, 1)$.

(The **Proof** *in Appendix 2.A* derives condition (2.31) in detail.) This lower bound on z is sufficient but not necessary: e.g., the analytic solution for the slope in Appendix 2.B.4 shows that $z \geq 0.2$ is sufficient when $g_0 = 2$ with $\lambda q\tau = 0.126$. Similarly, numerical solutions in Fig. 2.D.2b demonstrate an increasing profile on a broader z range: in the subpanels for z rising from 0.1 to 0.4, positive shocks to S^{bid} are too rare for the JE to prevail over the DPE. Once $z \geq 0.5$, bids create sufficient variance in S^{bid} for a positive slope. Increasing z further initially raises the slope but the slope falls back towards zero as z approaches 1 because zero-types are insensitive to S^{bid} changes and the F_{cc} and F_c terms contain $1 - z$.

The increasing profiles for $z \geq 0.5$ are mildly concave as the JE is stronger early on. More convex CDFs in the next subsection yield a stronger JE near the deadline.

2.4.3 Generic power distributions

The power-ρ CDF

$$F(c) = z + (1 - z)\left(\frac{c}{q}\right)^{\rho} \qquad \rho \geq 0 \tag{2.32}$$

generates bid profiles in Fig. 2.5 with slope moving from negative to positive as ρ varies from 1/2 to 3. The CDF is concave for $\rho \in (0, 1]$, explaining the decreasing bid profile at $\rho = 1/2$. We already discussed $\rho = 1$ and 2 so we now focus on $\rho > 2$. The bid dynamics are close to the quadratic case but the slope is more pronounced. Notice that S_0 is lower in this final panel because raising ρ places greater probability weight on high

(a) $\rho = 1/2$
($S_0 = 0.85$)

(b) $\rho = 1$
($S_0 = 0.72$)

(c) $\rho = 2$
($S_0 = 0.18$)

(d) $\rho = 3$
($S_0 = 0.17$)
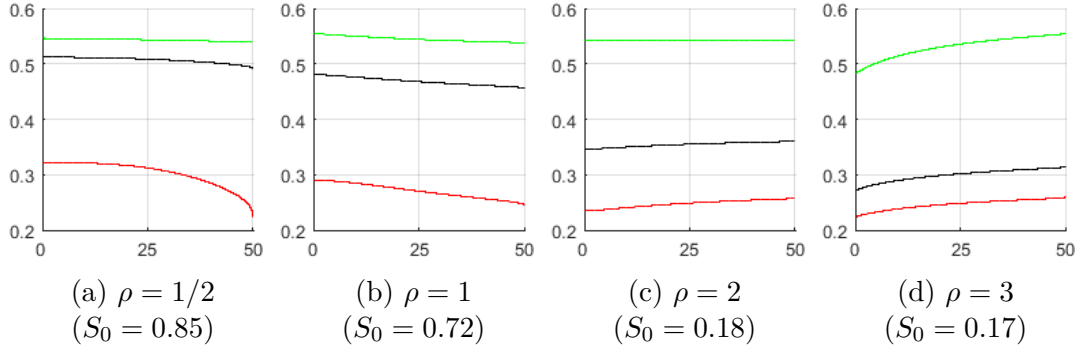
Figure 2.5: Average bids —— $A_t^{\mathcal{S}}$ —— A_t —— $A_t^{\mathcal{F}}$ against time t, for a power-ρ CDF with $z = 0.5$ and a campaign with $g_0{=}20$, $(\tau, \lambda, q){=}(50, 0.7, 0.75)$

inspection costs and Lemma 2.4 applies. This also affects the slope but the slope remains steeper when z is used to give similar success rate S_0: the profile for $z = 0.8, \rho = 3$ in the comparative statics Fig. 2.D.2c is steeper than that for $z = 0.7, \rho = 2$, with a similar S_0 in Fig. 2.D.2b.

2.4.4 Single-peaked distributions

In this subsection we consider single-peaked densities for the empirical distribution of the underlying net cost variable *prior* to its truncation. In such settings, the Jensen effect (JE) has a constant sign if the modal inspection cost is either low or high enough. In the low case, a negative mode implies that the truncated cost distribution has an atom at zero and decreasing cost density over the $(0, q]$ interval. That implies a concave CDF and a negative JE. In the high case, if the mode is bigger than q, the truncated cost distribution has a density that is increasing over $[0, q]$. That implies a convex CDF and generates a positive JE.

Fig. 2.6a illustrates a single-peaked probability density function (PDF) with mode $\mu_L \leq 0$ in brown, giving a concave CDF over $[0, q]$, while the high mode case $\mu_H \geq q$ in blue generates a convex CDF. Fig. 2.6b presents the concave and convex CDFs and Fig. 2.6c shows the corresponding bid profiles: decreasing for the distribution with $\mu = \mu_L = 0$ (intensity then falls from 0.35 to 0.25); increasing for the distribution with $\mu_H = 0.75$ (bid intensity increases from 0.3 to 0.36). In this illustration, $q = 0.75$ so it sets μ_L and μ_H at exactly 0 and q. In addition, the inspection costs follow a truncated normal distribution but we mix with an atom on zero-types so that the final atom of zero types is the same for both CDF's and to maintain comparable success rates, we compensate high μ_H with $\lambda_H > \lambda_L$ ($S_0 = 0.285$ for H and 0.250 for L). Variation in the zero atom can be interpreted as variation in the number of friends and contacts of the entrepreneur.

(a) PDF $\varphi(c; \mu)$ implying concave (μ_L) & convex (μ_H) CDF.♦

(b) Distribution functions
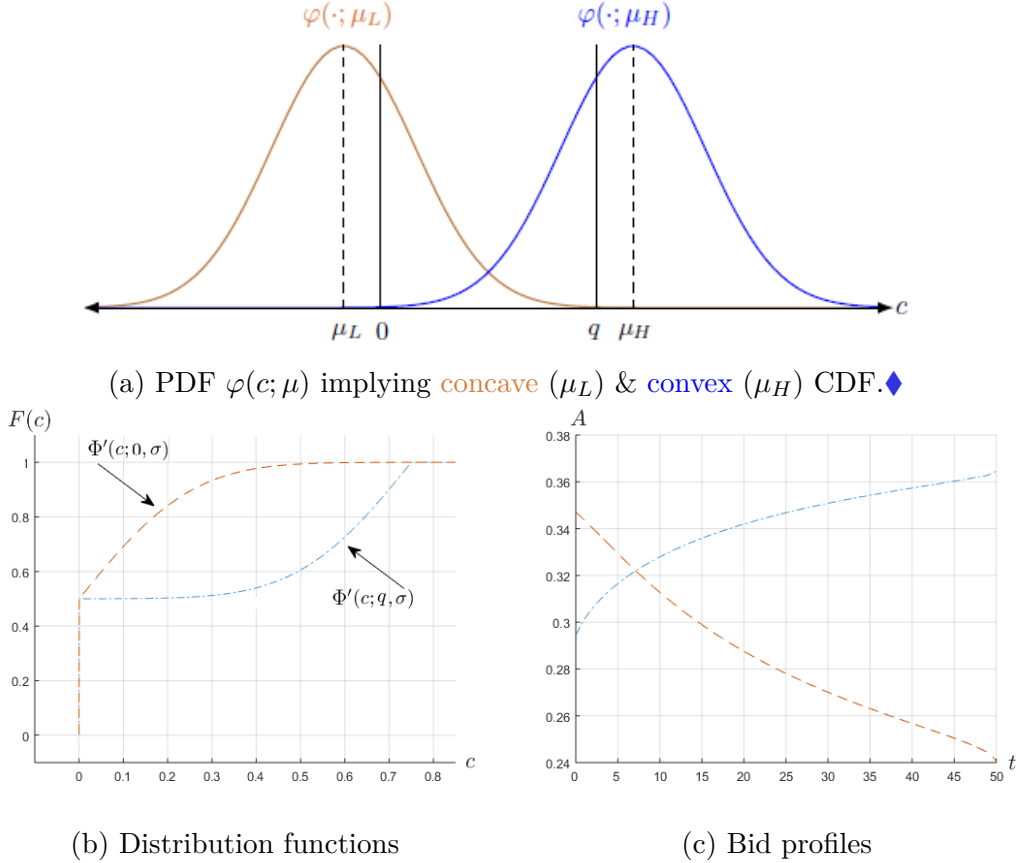
(c) Bid profiles

Figure 2.6: Truncated normal distributions $\Phi'(c; \mu, \sigma)$ on support $[0, q]$.   (a) Original PDFs and truncated region; (b) truncated CDFs; (c) bid profiles for $(\tau, q, z, \sigma) = (50, 0.75, 0.75, 0.5)$ and $(\lambda_L, \mu_L) = (0.5, 0)$ (concave $F = \Phi'(c; 0, \sigma)$), $(\lambda_H, \mu_H) = (0.75, 0.75)$ (convex $F = \Phi'(c; q, \sigma)$ ♦

2.4.5   The homogenous case

Turning to homogenous costs, we link bidding to survival probabilities. This allows us to succinctly characterize bid and success rates. We also demonstrate the discrete bid rate drops at the critical dates defined in Section 2.3.4 and show why the decreasing pivotality effect (DPE) dominates the Jensen effect (JE) at such dates and why both equal zero at all other dates.

When all bidders face the same inspection cost $c_t = c$, the CDF is a unit step function with jump discontinuity at c: $F(c_t) = 1$ if $c_t \geq c$ and 0 otherwise. If $c \leq 0$ or $c \geq q$, the average bid profile is flat because bids follow a homogenous Poisson process with intensity λq or 0, respectively. We therefore focus on the interesting case where $c \in (0, q)$.

Since bidders are now identical, every state (t, g) can be categorized as *active* if bidders choose $\mathbf{C}$ at (t, g) or *frozen* if they choose $\mathbf{A}$. The frozen state is absorbing because the gap cannot fall over an interval in which all bidders play $\mathbf{A}$. Moreover, time passing at any fixed gap g can only lower $S^{\text{bid}}_{(x,g)}$. So $S^{\text{bid}}_{(x,g)}$ stays at $S^{\text{bid}}_{(t,g)} = 0$ as x rises above

t. The frontier between the frozen and active states exactly corresponds to the critical dates introduced in Section 2.3.4. When $g \leq 1$, the project is active no matter how little time is left but for any gap $g \geq 2$, the campaign freezes when time remaining falls below the minimal duration at which cost c bidders are willing to bid which is given by $\hat{\tau}_g : S^{\text{bid}}_{(0,g)}(\hat{\tau}_g) = c/q$.[10] This is well-defined because $S^{\text{bid}}_{(t,g)}(\tau) \equiv S^{\text{bid}}_{(0,g)}(\tau - t)$ is continuous and decreasing in t and ranges over $[0,1)$: with enough time remaining, the success rate becomes arbitrarily close to one for any $\lambda q > 0$; conversely, for any $g \geq 2$, success after one bid is impossible when time runs out at $t = \tau$. To ease exposition, we mostly use associated critical dates $\hat{t}_g \triangleq \tau - \hat{\tau}_g$. In sum, for any $g \geq 2$

$$S^{\text{bid}}_{(\hat{t}_g, g)} \equiv S^{\text{bid}}_{(\tau - \hat{\tau}_g, g)}(\tau) \equiv S^{\text{bid}}_{(0,g)}(\hat{\tau}_g) \equiv \frac{c}{q} \tag{2.33}$$

The bid rate is λq while active, since any arriving bidder inspects and therefore bids with probability q. The atom size $z = 1$ and so the fall to zero on freezing corresponds to the DPE, $\mathcal{E}^{(t,g)}_{t_+} = -z = -1$ at $t = \hat{t}_g \triangleq \tau - \hat{\tau}_g$. The DPE is zero at all other dates since homogenous cost implies that DP affects all bidding decisions when it affects any.

The frontier between active and frozen states can also be described by the critical gap at which the campaign is just active for any given date t. We define the maximal gaps for activity as $\hat{\boldsymbol{g}} = \left(\hat{g}_t\right)_{t \in [0,\tau]}$ where

$$\hat{g}_t \triangleq \sup \left\{ g \in \mathbb{Z} : S^{\text{bid}}_{(t,g)} \geq \frac{c}{q} \right\} \tag{2.34}$$

These critical dates and gaps trace out the frontier. We call it a *wall of ice*, since a campaign instantly freezes when its path crosses into the region with $g_t > \hat{g}_t$. A campaign is active at t if its trajectory $\boldsymbol{g} = \left(g_t\right)_{t \in [0,\tau]}$ has not crossed $\hat{\boldsymbol{g}}$ before t. We illustrate this wall of ice in violet in Fig. 2.7(a); it separates active (below) from frozen regions (above).

Setting $\hat{g}_{\tau_+} \triangleq 0$, success is equivalent to staying weakly below the wall until τ_+. If $g_0 > \hat{g}_0$, equivalent to $\tau < \hat{\tau}_{g_0}$ or $\hat{t}_{g_0} < 0$, the campaign is born frozen so we focus on $g_0 \leq \hat{g}_0$. Fig. 2.7(a) also exhibits four specific gap paths. Paths (1) and (2) are in red since they end up failing: they cross the violet wall of ice at $t = 23.5$ and 49.5 where $\hat{g}_{23.5} = 14$ and $\hat{g}_{49.5} = 2$, respectively. From then on, they are necessarily flat. (3) and (4), in green, successfully stay below the wall of ice; $g_t \leq \hat{\boldsymbol{g}}_t$ for all $t \leq \tau$ and $g_{\tau_+} \leq \hat{g}_{\tau_+} \equiv 0$.

Decreasing bid profile. Since the frozen state is absorbing, its probability can only increase over time. Bidding intensity is λq when active and zero when frozen. This proves that the average bid profile is decreasing. Fig. 2.7(b) illustrates alongside the parallel profile for pivotality of Fig. 2.7(c).

Average bidding falls via discrete drops at precisely the same critical dates as the vertical drops in the violet wall of ice in Fig. 2.7(a). Notice that the magnitudes of

---

[10]For $g \leq 1$, $\hat{\tau}_g = 0$ since $S^{\text{bid}}_{(t,g)} = 1, \forall t \leq \tau$. With only one atom, $K = 1$, $z^0 = 0$ and we drop index k.

(a) Wall of ice   (b) Bid profiles   (c) Pivotality profile
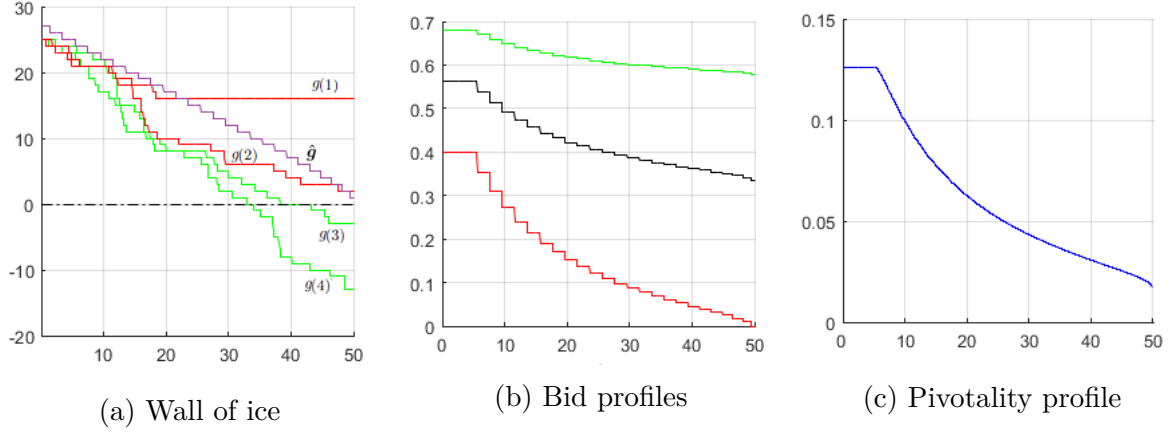
Figure 2.7: Wall of ice, four simulated gap paths and profiles of bids and pivotality for homogeneous inspection cost c=0.2, g_0=25, (τ,λ,q)=(50,0.75,0.75), giving $S_0 = 0.58$

Legend: (a) — $\hat{g}_t$ — $g(\bullet)$ successful — $g(\bullet)$ failed   (b) — $A_t^{\mathcal{S}}$ — A_t — $A_t^{\mathcal{F}}$ (c) — $A_t^{\Delta S}$

the drops are decreasing. Each drop equals λq times the unit atom from homogeneity times the probability of hitting the vertical wall. This probability of freezing equals the probability that the gap is critical at the critical date, $g_{\hat{t}_g} = g$. It falls over time because paths that cross the wall by t never hit it again, while paths that diverge below the wall rarely come back to cross it. Since the wall-of-ice is approximately linear, these falling jump sizes make the bid profile approximately convex from the first critical date onwards.

The initial, perfectly flat bidding plateau owes to the fact that g_0 must start below the critical gap in a campaign that is not born frozen. Given $g_0 < \hat{g}_0(\tau)$, the campaign faces no risk of freezing until $t = \hat{t}_{g_0}$. Larger durations τ extend this initial plateau and diminish the size of the downward steps along with the probability that the campaign ever freezes, the probability of a failure.

Decomposing into DP and JE adds little insight in this special case but will prove useful when we move to multiple atoms and when we combine continuous and discrete distributions. As noted above, the DPE is zero except at critical dates because marginal changes in $S_{(t,g)}^{\text{bid}}$ only affect bidding $\beta_{(t,g)} = \lambda q F\left(q S_{(t,g)}^{\text{bid}}\right)$ at dates when $q S_{(t,g)}^{\text{bid}} = c$; this is immediate from Eq. (2.15) which applies at all non-critical dates since $F(\cdot)$ is only discontinuous at c. The JE at non-critical dates is given by Eq. (2.17) as $\mathcal{J}_{(t,g)} = \beta_{(t,g)}\Delta F(q S_{(t,g)}^{\text{bid}})$ since $F_c = 0$ at non-critical dates.

This product always equals zero in the homogenous case because the difference term is zero for gaps below the wall of ice and the bid rate is zero for gaps strictly above the wall of ice. The logic resides in the JE's two necessary components. First, local convexity or concavity here requires a non-zero difference ΔF but ΔF is zero below the wall of ice because gap reduction cannot raise the inspection probability when already maximized at unity. Strictly prior to a critical date, the unique cost type always inspects even if a

small amount of time passes with no bid. That is, the project is maximally and hence equally active after both good and bad news. Second, uncertainty requires the possibility of a bid but strictly above the frontier the campaign is frozen and generates no news at all so the JE is zero there too. For a quick mathematical proof: note that $\Delta F(qS^{\text{bid}}_{(t,g)}) = 0$ except when the gap g_t is one unit above the critical gap for that high cost; at such gaps, $\Delta F(qS^{\text{bid}}_{(t,g)}) = 1$ but then $\beta_{(t,g)} = 0$ at all such gaps and $\mathcal{J}^{(t,g)} = 0$ even there.

Note that average pivotality decreases smoothly but the profile in Fig. 2.7(c) is flat until the first critical date. The reason why DP is trivial here is that on interval $[0, \hat{t}_{g_0}]$ any arriving bidder inspects. The earlier of two bidders who may arrive on $[0, \hat{t}_{g_0}]$ has more successors but cannot influence them since they are already maximally active. So the earlier bidder has no encouragement effect on the extra successors. There is no scope for additional strategic complementarity, fixing average pivotality until $\hat{t}_{g_0}$. The reason why average pivotality profile of (c) is continuous is that success rates depend on bidding over time and are given by integration of a finite function over time, shown in Eq. (REC-S). It is the step discontinuity in $F(\cdot)$ that generates the temporal discontinuities in bidding $\beta_{(t,g_t)}$ which depends on $F(S_{(t,g_t)} + \Delta S_{(t,g_t)})$ by Eq. (2.6).

Turning to critical dates, Fig. 2.8a illustrates how the DPE's discrete negative effect always dominates the JE at such moments and creates a discrete downward jump. Fig. 2.8a, like Fig. 2.2, considers a small time interval ϵ now starting from $t = \hat{t}_g$, the critical date corresponding to a generic gap g. The DP arrow shown in magenta is negative. It is only of order ϵ, but has a discrete effect shown by the DPE in green because $F(\cdot)$ is discontinuous at t: $\mathcal{E}^{(t,g)}_{t+\epsilon}$ equals -1 for any $\epsilon > 0$. Uncertainty in $S^{\text{bid}}_{(t+\epsilon,\tilde{g}_{t+\epsilon})}$ around $\mathbb{E}S^{\text{bid}}_{(t+\epsilon,\tilde{g}_{t+\epsilon})}$ creates the positive JE $\mathcal{J}^{(t,g)}_{t+\epsilon}$ shown in orange but its magnitude is of order ϵ. So it shrinks to zero and is dominated by the negative DPE as $\epsilon \to 0_+$. In brief, the JE is positive because $\mathbb{E}S^{\text{bid}}_{(t+\epsilon,\tilde{g}_{t+\epsilon})}$ lies just below c/q but JE is continuous so the jump is infinitesimal; to understand why the JE is positive in terms of the frontier, notice that $t = \hat{t}^c_g$ in the figure, $\hat{g}^c_t = g$ and $\hat{g}^c_t = g - 1$ so at $t + \epsilon$, the gap g is exactly one unit above the frontier. The key, infinitesimal result follows from the fact that the chance of a bid arriving in any instant is zero. By contrast, the DPE is discrete. The project instantly freezes if no bid arrives at the critical date. Formally, as ϵ goes to zero, $S^{\text{bid}}_{(t+\epsilon,g)} \to S^{\text{bid}}_{(t,g)}$ so that DP $\mathcal{D} \to 0$ (and JE $\mathcal{J} \to 0$) but DPE $\mathcal{E}$ stays fixed at -1.

Finally, we discuss the fanning out, anticipated in Section 2.3.6, of the conditional curves in Fig. 2.7(b). This figure presents the profiles for bid averages conditioned on success (in green) and failure (in red). Qualitatively similar to the unconditional average profile, the downward steps in the red profile are larger because conditioning on failure raises the probability of freezing during the campaign (to one minus the probability of failing with $g = 1$). It might come as a surprise that the success-conditional profile has any downward steps because success conditioning precludes crossing the wall of ice. However, the wall still matters: successful paths must have enough early bids to stay

(a) Homogeneous inspection cost ♦

(b) Binary inspection cost (c, c_L) with $c > c_L = 0$ and atom z^0 on c_L ♦

Figure 2.8: DP, DPE and JE for discrete distributions

below the wall. As explained in Section 2.3.6, success conditioning implies a selective bias towards bidding rates above λq. This bid-selection bias is highest early on because survival requires n bids by the critical date $\hat{t}_{g_0-n}$ for each $n \geq 1$. That is why the green curve has downward steps at critical dates.

Failure-conditioned paths have a negative bid-selection bias so the profile is always lower. A failure can have no bids but *given* any positive aggregate number of bids below g_0, the bid arrivals must again be sufficiently early for the project to survive until the last of those bids is placed. This again creates downward steps. The downsteps are larger than in the green success-conditional profile because bidding paths that end up failing are usually closer to the wall of ice, giving it a greater influence. In the terms of Section 2.3.6, selective weighting on lower gaps when conditioning on failure again has no impact at the very start of a campaign but kicks in after the first critical date and contributes to the larger size of the downward steps in the red profile compared to the green profile.

Succinct characterization of bidding and success rates.

The active-frozen dichotomy in the homogenous case permits an explicit characterization of bid rates via survival probabilities. We replace Proposition 2.1's generic integral-based recursion for success rates with a finite recursive sum. Section 2.3.6's transition probabilities give the probability that a campaign survives from state (t, g) till $t' > t$:

$$\alpha_{t'}^{(t,g)} \triangleq \sum_{g' \leq \hat{g}_{t'}} Q_{(t',g')}^{(t,g)} \tag{2.35}$$

The probability of surviving till t given starting gap g_0 is $\alpha_t^{(0,g_0)}$. The average bid rate is $A_t = \lambda q \alpha_t^{(0,g_0)}$ and the success rate is $S_{(t,g)} = \alpha_{\tau_+}^{(t,g)}$. Defining Poisson probability function,

$$\mathcal{P}(b; \Lambda) \triangleq \frac{\left(\Lambda\right)^b e^{-\Lambda}}{b!} \tag{2.36}$$

for $b \geq 0$ bidding events and Poisson parameter $\Lambda \geq 0$, we prove

Proposition 2.7. *Success and average bid rates under homogeneity are characterized by*

$$S_{(t,g)} = \alpha_{\tau_+}^{(t,g)} \qquad A_t = \lambda q \alpha_t^{(0,g_0)} \tag{2.37}$$

where $\forall\ g, t \leq t'$, (i) $\alpha_{t'}^{(t,g)} = 1$ if $t' \leq \hat{t}_g$, (ii) $\alpha_{t'}^{(t,g)} = 0$ if $\hat{t}_g < t$, (iii) on $t \leq \hat{t}_g < t' \leq \tau_+$,

$$1 - \alpha_{t'}^{(t,g)} = \mathcal{P}\left(0; \lambda q(\hat{t}_g - t)\right) + \sum_{b=1}^{g-\hat{g}_{t'}-1} \mathcal{P}\left(b; \lambda q(\hat{t}_g - t)\right)\left(1 - \alpha_{t'}^{(\hat{t}_g, g-b)}\right) \tag{REC-α}$$

$$1 - c/q = \mathcal{P}\left(0; \lambda q(\hat{t}_{g-1} - \hat{t}_g)\right) + \sum_{b=1}^{g-2} \mathcal{P}\left(b; \lambda q(\hat{t}_{g-1} - \hat{t}_g)\right)\left(1 - \alpha_{\tau_+}^{(\hat{t}_{g-1}, g-1-b)}\right) \tag{REC-$\hat{t}_g$}$$

Remark. Before proving, we state the immediate corollary recursion (REC-S-hom) to

find $S_{(t,g)}$ given $S_{(\hat{t}_g,g-1)}, ..., S_{(\hat{t}_g,1)}$ and $\hat{t}_g$; *hom* indicates homogeneity. Since (REC-$\hat{t}_g$) can be solved for $\hat{t}_g$ given $\hat{t}_{g-1}$ and $S_{(\hat{t}_g,g-1)}, ..., S_{(\hat{t}_g,1)}$, we can combine these recursions to solve for both $S_{(t,g)}$ and $\hat{t}_g$ given their solutions at gaps $g-1$ and below; recall $\hat{t}_1 = \tau$.

Corollary. $S_{(t,g)} \equiv 1$ for all $g \leq 0$ initiates a recursive solution for generic $S_{(t,g)}$ via

$$1 - S_{(t,g)} = \mathcal{P}\left(0; \lambda q(\hat{t}_g - t)\right) + \sum_{b=1}^{g-1} \mathcal{P}\left(b; \lambda q(\hat{t}_g - t)\right)\left(1 - S_{(\hat{t}_g, g-b)}\right) \qquad \text{(REC-}S\text{-hom)}$$

Proof. (i) If $t' \leq \hat{t}_g$, $\alpha_{t'}^{(t,g)} = 1$ since the project is safely below the wall of ice and cannot freeze between t and t'. (ii) If $t > \hat{t}_g$, the campaign is already frozen at t so it can neither survive nor succeed: $\alpha_{t'}^{(t,g)} = \alpha_{\tau_+}^{(t,g)} = 0$. (iii) If $t \leq \hat{t}_g < t'$, both freezing and survival are possible. On $[t, \hat{t}_g]$, bids arrive with homogenous intensity λq so b has Poisson parameter $\Lambda = \lambda q\left(\hat{t}_g - t\right)$ on that interval. The probability of failing to survive till t' is the sum of the probability that $b = 0$ or, some $b \in \{1, ..., g - \hat{g}_{t'} - 1\}$ bids arrive by $\hat{t}_g$ *and* then the campaign fails to survive from state $\left(\hat{t}_g, g - b\right)$ till t', proving Eq. (REC-α). Notice that any $g \leq 0$ guarantees success and implies $\hat{t}_g = \tau_+$ and $\alpha_{\tau_+}^{(t,g)} = 1$ for any $t \leq \tau_+$; this provides the initial step in Eq. (REC-α).

Recursion (REC-$\hat{t}_g$) for $\hat{t}_g$ at $g \geq 1$ combines $S_{(\hat{t}_g, g-1)} = S_{(\hat{t}_g, g)}^{\text{bid}} = c/q$ from Eq. (2.33) with recursion Eq. (REC-α) at $t' = \tau_+$ and starting state $(\hat{t}_g, g-1)$; recall that $\hat{g}_{\tau_+} = 0$. The initial step for $\hat{t}_g$ at $g = 1$ follows from the general result that $\hat{t}_1 = \tau$ since at $g = 1$ bidders play **C** at any date during the campaign; conversely, $\hat{g}_\tau = 1$. ∎

Proposition 2.7's explicit linear recursion speeds up computations more than ten-fold which is useful for optimizing design since numerical calculations become more intensive there. The fact that the violet curve in Fig. 2.7(a) is approximately linear when g and τ get large is also useful for design because it suggests that the maximal effective thresholds increase approximately linearly with the expected number of bidders in the homogenous setting.

Illustration with initial gap $g_0 = 2$. $S_0 = S_{(0,2)} = 1 - \left[e^{-\lambda q \hat{t}_2} + \left(\lambda q \hat{t}_2 e^{-\lambda q \hat{t}_2}\right) e^{-\lambda q(\tau - \hat{t}_2)}\right]$; this is the probability of collecting at least two bids, with at least one by $\hat{t}_2$. The first term in the bracket is the probability of no bids on $[0, \hat{t}_2]$; the second term is the probability of one on $[0, \hat{t}_2]$ but no bid on $[\hat{t}_2, \tau]$ which lasts $\tau - \hat{t}_2$. Simplifying, $S_0 = 1 - \left[e^{-\lambda q \hat{t}_2} + \lambda q \hat{t}_2 e^{-\lambda q \tau}\right]$.

While $g = 2$, bidders only inspect if $t \leq \hat{t}_2$, whereas once $g \leq 1$, bidders always inspect. So the average bid rate is constant at $A_t = \lambda q$ up till $\hat{t}_2$ and drops at that date by λq times the probability $1 - \alpha_{\hat{t}_2}^{(0,2)} = Q_{(\hat{t}_2,2)}^{(0,2)} = e^{-\lambda q \hat{t}_2}$ of hitting the vertical wall of ice at $\hat{t}_2$. It then remains constant at $A_t = \lambda q\left(1 - e^{-\lambda q \hat{t}_2}\right)$. $S_{(\hat{t}_2,1)} = 1 - e^{-\lambda q(\tau - \hat{t}_2)} = S_{(\hat{t}_2,1)} = c/q$, so $\hat{t}_2 = \tau - \frac{1}{\lambda q} \ln\left(1 - \frac{c}{q}\right)^{-1}$ or more simply, $e^{-\lambda q \hat{t}_2} = 1 - c/q$. Intuitively, $\hat{t}_2$ is lower in

adverse settings because bidders at $g = 2$ then give up earlier: $\hat{t}_2$ is decreasing in c and increasing in λ and q.

2.4.6 Richer discrete distributions

In this section, we first show how multiple atoms lead to positive Jensen effects (JE) associated with each atom except for the lowest atom.[11] In the binary case, decreasing pivotality effects (DPE) are again null except at the critical dates corresponding to each of the two types. Combined with the continuous positive JE associated with the higher type, this creates a tooth-shaped profile. We explain the exact origin of the positive JEs and how they combine additively, weighted by the distribution of the gap at any given date.

We illustrate with three binary examples. In the first Fig. 2.9a, there is a zero atom ($c_L = 0$) and so the only DPE is that associated with $c_H > 0$. In addition, the high cost c_H is moderate which results in moderately strong JEs as soon as the first critical date has passed. In that case the JE's are decreasing over time because the rate of increase in entropy in S^{bid} decreases over time, on average.

In the second Fig. 2.9b, there is again a zero atom ($c_L = 0$) but now the high cost is extremely high, $c_H = q$, the upper limit value. As a result, the project starts in the cold state where only zero types inspect and the DPE is trivial. DPE is trivial because the critical date for zero types is after the deadline and the frontier for the high cost $c_H = q$ is perfectly flat at $\hat{g}_t^H \equiv 1$, so it is never crossed from below. Once H-types are willing to inspect, they are always willing to inspect. That is, the hot state is an absorbing state in this special case. As a result, the profile is upward sloping. We highlight also that the upward slope is most prominent in the later phases of the campaign when the atom is at a high cost equal to or close to the upper limit value q.

In the third Fig. 2.9c, there is no atom at 0. As a result, while the project starts in the warm state where the L-types inspect, giving bid rate $\lambda q z_L = 0.6$, the DPE generates a flight of downward steps that is most pronounced early on. This early downward slope is very similar to that from the homogenous case. The reason for this is that the H-type again has the limit cost $c_H = q$ so the possibility of the project getting hot in that even H-types inspect is very remote in the early phase and has negligible impact on the profile. In the second half of the campaign, by contrast, the possibility of heating up and activating even the limit type is increasingly relevant. The JE associated with the H-type generates the positive slope towards the end of the campaign. This generates a clear U-shape. It is especially pronounced when conditioning on success because the probability that the limit or H-type is activated on a failing project is very remote.

---

[11]The next subsection shows that, when combined with a continuous distribution, the lowest atom also generates a positive JE if that atom has cost $c > 0$.
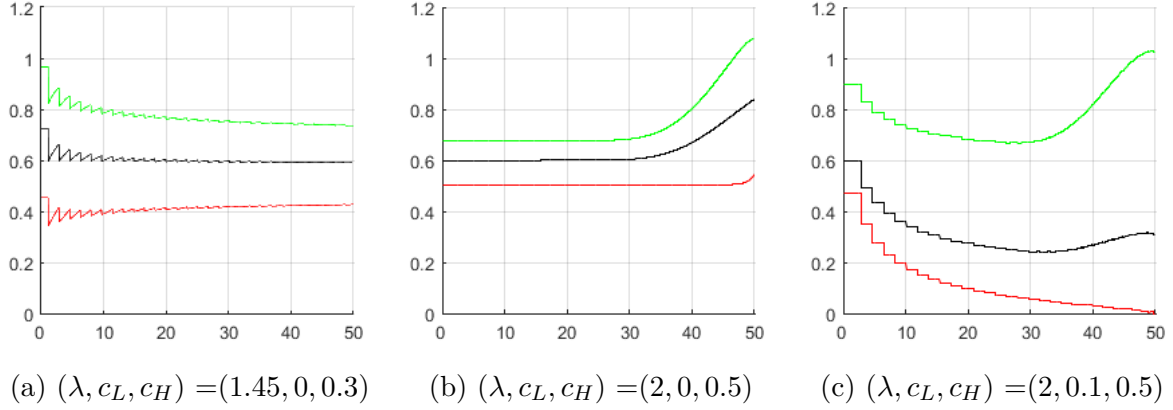
(a) $(\lambda, c_L, c_H) =(1.45, 0, 0.3)$ (b) $(\lambda, c_L, c_H) =(2, 0, 0.5)$ (c) $(\lambda, c_L, c_H) =(2, 0.1, 0.5)$

Figure 2.9: Bid profiles for binary inspection costs, $c_t \in \{c_L, c_H\}$ with probability $z = 0.6$ on c_L, $1 - z = 0.4$ on c_H, for $g_0 = 30$, $(\tau, q) = (50, 0.5)$

Turning to richer discrete distributions, the lowest cost atom, if non-zero, defines a wall-of-ice exactly as in the homogenous case but in contrast to that case, if a cost c^j atom is not the lowest atom, then $F(c^j) - z^j > 0$ which implies that the campaign slows down but does not immediately freeze in states lying above the critical date-gap frontier corresponding to atom j. From states just above this frontier, good news activates this atom z^j of bidders and this creates a positive JE. Since the DPE is zero except at critical dates, the net effect of time is positive, creating non-monotonic saw-tooth profiles.

The discontinuity exhibited by the bid profile under homogeneity of inspection costs appear also with distributions of inspection costs over a discrete set of values c^k, with $c^1 < c^2 < \ldots, < c^K$ such as the two-point distribution on $c^1 \equiv c^L, c^2 \equiv c^H > c^L$. Each c_k determines a contour set $\hat{g}^k_t$ for k-type bidders. This translates into at most $K + 1$ bidding states: states $1, \ldots, K$ such that all bidders with $k' \leq k$ inspect and, if present, an additional frozen state 0 where no bidder inspects. The typical profile generated by these distributions is tooth-shaped with discontinuity points at each $\hat{t}^k_g \in [0, \tau]$.

Fig. 2.9 illustrates for the two-type distribution with $c^L = 0$ and $z = 0.6$. As in earlier sections, $z \in (0, 1)$ denotes the probability on the zero-type c^L. Since zero-types always inspect, there is no frozen state and the set of critical dates are just those corresponding to the c^H type. These critical dates and the corresponding critical gaps now map out a hot-cold boundary. As with the wall of ice, it again consists of horizontal and vertical segments, but now the probability than an arrival inspects drops from 1 to z when the project moves from the hot into the cold state (by crossing a vertical segment at a critical date) so the project may move back into the hot state later on. Entering or re-entering the hot state happens when a bid causes the gap to fall down to $\hat{g}^H_t$. At such moments, (t, g) hits a horizontal segment of the hot-cold boundary and the bidding rate accelerates from $\lambda z q > 0$ to λq. Acceleration events can occur at any dates at which the gap may strictly exceed $\hat{g}^H_t$. That is, at any date after $\hat{t}^H_{g_t}$. These acceleration possibilities explain why the bidding profile has an upward slope between critical dates. There is a positive

Jensen effect.

This can be understood from Fig. 2.8b. Now $S^{\text{bid}}_{(t,g)} = S^{\text{bid}}_{(t,\hat{g}_t+1)}$ is slightly to the left of the critical value c/q so that $F(S^{\text{bid}}_{(t,g)}) = z^0$; as before, $F(S^{\text{bid}}_{(t+\epsilon,\hat{g}_t)}) = 1$ after the good news of a bid indicated by the upper triangle. The DPE arrow is invisible since it is null. The upward orange arrow indicates that the positive JE is the only effect and this generates the positive slope.

We also provide a non-graphical explanation. Notice that for any t, there is a positive probability that $g_t = \hat{g}_t + 1$. Then Eq. (2.20) adjusts to $\mathcal{L}^{\beta}_{(t,g)} = \lambda q \beta_{(t,g)} \Delta F$ with $\Delta F = 1 - z^0 > 0$ and $\beta_{(t,g)} = \lambda q z^0$. So $\mathcal{L}^{\beta}_{(t,g)}(\lambda q)^2 z^0(1 - z^0)$ is also positive. This JE has a fixed magnitude and is positive only in states $(t, \hat{g}_t+1)$. The probability of reaching such states sometimes falls and sometimes rises as critical dates get near, creating slight concavities and convexities in the upward slopes of the teeth depicted in Fig. 2.9a.

In general, when F consists of atoms, away from critical dates, the DPE is zero and the JE is always positive. This is because one moment after a generic date, in the bad news event of no bid, the marginal falls in S^{bid} which is time continuous have no effect on $F(qS^{\text{bid}})$, whereas the good news event of a bid has a discrete positive effect which shifts the bid rate upwards whenever the gap is one above the frontier for any cost atom.

Time's overall average rate of impact is given by Eq. (2.19) as $\lambda q \left(\mathcal{J}_{(t,g)} + \mathcal{E}_{(t,g)}\right)$. For any (t, g), there is at most one bidder type k for which $t = \hat{t}^k_g$ and generically there are none. At those critical dates, $\mathcal{E}^{(t,g)} = -\infty$ i.e. discrete drop $-z^k$ where $z^k = \mathbb{P}\left(c = c^k\right)$. At all other dates $\mathcal{E}_{(t,g)} = 0$.

Jensen effects occur over a wider range of possible dates and multiple types may start to inspect after a unit fall in the gap. Nonetheless, for any type j and date t, there is at most one such gap $g = \hat{g}^j_t + 1$ at which the type j JE is positive.

$$\mathcal{J}_{(t,g)} = \beta_{(t,g)} \sum_{k:g=\hat{g}^k_t+1} z^k$$

To predict the profile, we use the fact that at non-critical dates, $\dot{A}_t = \sum_{g=-\infty}^{g_0} Q_{(t,g)}\mathcal{L}^{\beta}_{(t,g)}$ (from Eq. (2.26)) and $\mathcal{L}^{\beta}_{(t,g)} = \lambda q \mathcal{J}_{(t,g)}$ given that the DPE is null. Meanwhile, at critical dates, the discrete drops in A_t are given by weighted sums of $z^k \mathbb{P}\left(g_t = \hat{g}^k_t\right)$.

2.4.7 The U-shape

This section explains how simple combinations of the above dynamics readily generate a U-shaped profile. We already achieved a rudimentary U-shape in the binary case but we now derive a smooth U-shape by combining discrete and continuous distributions in a less rigid setup.

First, note that the classic uniform cost distribution already provides a good fit for the initial downward slope of the U-shape. As shown in Section 2.4.1 and illustrated well

by Fig. 2.4, the decreasing pivotality effect (DPE) is strongest at the beginning because pivotality usually gets close to zero as a campaign matures. Second, the strong positive Jensen effect (JE) associated with a substantial subgroup of bidders with high inspection costs is largely confined to near the end of the campaign, as shown in panels (b) and (c) of Fig. 2.9. The reason is that the JE is trivial until the gap is very low, concretely at $g = 2$ in the example corresponding to Fig. 2.9. Such low gaps are usually only reached towards the end of a campaign. For that reason, the JE generates a final upward slope and creates a U-shape. Both the logic behind a strong DPE early in a campaign and the logic behind a strong positive JE late in a campaign are robust to combining these two distributions as the next example demonstrates.

A smoother variant that also generates a U-shaped profile is that from a cost distribution that combines the flat density function of the uniform distribution or a decreasing density, with an increasing density such as for the quadratic CDF of Section 2.4.2 or higher power CDFs. If we take a linear decreasing and then linear increasing density, we essentially derive a U-shaped bid profile from a V-shaped density. The question is why any such distribution might be plausible.

Consider the case of a relatively flat distribution such as the uniform distribution combined with an atom on $c = q$ as in Figs. 2.9b and 2.9c. As explained in the introduction to this section, our argument is that bidders can be split into two basic groups: those in group 1 take account of their inspection costs as in the model; those in group 2 only bother to consider sure options, which correspond here sales offers that the bidder is guaranteed to be able to buy if he wants it. These are the rule-of-thumb bidders we described above. Recall that these group 2 bidders are equivalent to bidders with cost $c = q$ because a type with the limit-cost, a cost on the verge of being prohibitive, will only inspect when $g \leq 1$ which is precisely where the sales offer to a given bidder becomes a sure offer, not contingent on what any later bidders will do.

We now formally model the group 2 actors. We assume an atom of size z^q at $c = q$ because they inspect when S^{bid} is near 1 that if g_0 is large enough, only occurs towards the end of the campaign. So we assume that the CDF is

$$F(c) = (1 - z^q)\frac{c}{q} \text{ for } c < q \text{ and } F(c) = 1 \text{ for } c = q$$

Clearly the q-type bidders start inspecting at the first t such that $g_t = 1$ and continue afterwards, i.e. $\hat{g}_t^q = 1$ for all t. This distinguishes two bidding phases, based on whether q-types are active (willing to inspect) or inactive and we say that the campaign is q-active if they are active. Extending the notation from Section 2.4.5, we denote the probability that the campaign is q-active at t given its current state (t', g') by $\alpha_t^{q(t',g')} = \sum_{g \leq 1} Q_{(t,g)}^{(t',g')}$.

Since the probability to switch into the q-active phase is time increasing, it generates a positive JE near the end of the campaign. On the other hand, the initial DPE or the
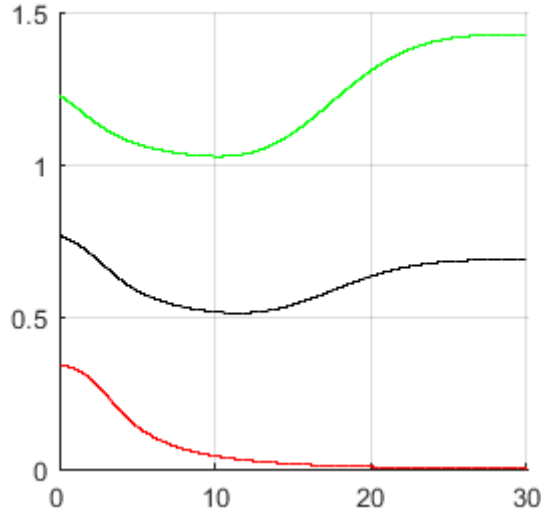
Figure 2.10: Bid profiles for uniform distribution with atom $z^R = 0.3$ on $c_R = q$ with $g_0 = 20$, $(\tau, \lambda, q) = (30, 1.5, 0.95)$

regular bidders is stronger early on, so the combined effect originates a U-shape. Both effects are present for successful campaigns, but the JE is excluded for failing campaigns since they must have a zero bid rate once q-active in order to fail. So their profile is decreasing.

Proposition 2.8. *A uniform inspection cost on support $[0, q]$ plus an atom z^R on the limit cost type $c_R = q$ generate single-troughed or U-shaped average and success-conditional bidding profiles. The failure-conditional bid profile is strictly decreasing.*

The bid profiles presented in Fig. 2.10 present the features described by the previous proposition. The average and success conditional profiles show a lull in bidding starting roughly at $t = 5$ and terminating at $t = 20$ when the bidding intensity becomes high again thanks to the q-type JE. The average profile shows an initial concave region that disappears in the success-conditional profile due to the success effect raising bidding early. The magnitude JE depends directly on the probability of a switch from cold to hot phase, so it increases rapidly when the switch becomes first possible and then reduces when it becomes very likely that the campaign is hot. The fail-conditional profile is decreasing since bids never occur when the campaign is hot.

2.5 Design

Our baseline model treats all campaign parameters as given. In reality, entrepreneurs and crowdfunding platforms directly choose some parameters and have some influence over others. This section uses our framework to investigate how entrepreneurs and platforms can optimize their design choices. We focus on the case of an entrepreneur who maximizes

her project's success rate subject to a funding constraint. She needs G units of money to fund the fixed costs of investing in production that she cannot cover with her own money and external credit. We show that, to maximize success, she always sets her funding threshold at G because a funding threshold below G precludes production, while a higher threshold lowers success rates for any price rule (see Corollary 2.1). We call this funding threshold her "goal", consistent with platforms like Kickstarter. So her optimal *funding* threshold or goal is G throughout this section but the *bidder* threshold g_0 depends on pricing. That bidder threshold g_0 and pricing p are now endogenous. We continue to leave bidder arrival intensity λ, product quality variables v, q and campaign duration τ as exogenous parameters but we do discuss those extensions in Section 2.6.

We first consider the single-reward setting of the baseline model. With only one price p, the entrepreneur's decision problem reduces to that of optimizing p given its impact on g_0 and bidder incentives to inspect. We illustrate an interior pricing solution when costs are uniformly distributed.

Then we expand to show how multiple prices can add value even though bidders are ex-ante identical. We do not let prices depend on time, consistent with most platforms, but we do analyze reward menus with limited quantities. The limited quantities result in a form of gap-dependent pricing; recall that there is only one possible product. The most general such pricing scheme specifies a different price for the product at each gap value. This has theoretical interest to understand pricing factors but realism suggests fewer prices. So to be more in line with practice, we focus on the restriction to a two tier pricing scheme.

2.5.1 Endogenous threshold and pricing with a single reward

This section covers the single price setting. The fixed price p is no longer equal to $v - 1$ so the associated discount d appears as $v - p$ instead of being invisible as a factor 1 in the inspection incentive or expected return to inspecting $dqS_{(t,g)}^{\text{bid}}$ which was written as $qS_{(t,g)}^{\text{bid}}$ in the baseline setting where $d \equiv 1$. The bidding intensity expression Eq. (2.6) now becomes

$$\beta_{(t,g)}(p) = \lambda q F\left((v - p)q S_{(t,g-1)}(p)\right)$$

Lemma 2.5 readily generalizes: for any state (t, g), any change that increases $\beta_{(t,g)}$ while fixing the value of $S_{(t,g-1)}$, also increases $S_{(t,g)}$ by the same induction logic used to prove Lemma 2.5. In particular, increasing d by reducing the bid price p raises $S_{(t,g)}$ for all g.

Formally, the entrepreneur's problem is to set g_0 and price p to maximize her success rate subject to the constraint of satisfying her funding needs or goal G:

$$\max_{g_0, p} S_0(p) \qquad \text{subject to} \qquad G = pg_0 \tag{2.38}$$

$S_0(p) \equiv S_{(0,g_0)}(p)$ is given by the natural generalization of the recursive solution in Proposition 2.1 provided that we can continue to neglect blind bidding **B** which we assume is valid and justify later. That is, we assume it is optimal to continue to satisfy the No Blind Bidding constraint which adjusts from Assumption 2.1 to

Constraint 1 (Single price NBB). $p > qv$

In solving for $S_{(t,g)}(p)$ given this constraint, the only direct effect from generalizing d from 1 to $v - p$ lies in the change to Eq. (2.6). We omit the formal restatement of the recursive solution here because it is straightforward and because it can be seen by restricting all prices to the same value in the solution for the general multiple pricing case which the next subsection presents in full.

The entrepreneur faces the following tradeoff in her choice of price p: lowering p makes inspecting and bidding more attractive given a fixed initial gap g_0 (because of strategic complementarity, $S_{(t,g)}(p)$ rises for all (t, g)), but getting less revenue per bidder obliges her to raise g_0 (she needs more bidders to bid in order to reach her financial goal of G) and this has a negative effect on $S_{(0,g_0)}$. So a success-maximizer only raises the price in order to lower $g_0 = \lceil G/p \rceil$. It follows that G/p is an integer and $g_0 = G/p$.

Lemma 2.7. *Success-maximization requires $p = G/g_0$.*

So the optimal threshold can be found by optimizing over integers g_0, setting $p = p(g_0) \triangleq G/g_0$ and imposing $g_0 < G/q$ since $p \leq v$ else nobody ever bids.[12] $p \leq v$ requires $g_0 \geq \lceil \frac{G}{v} \rceil$. We normalize v to 1 and therefore only need consider gaps $g_0 \geq \lceil G \rceil$. Notice that the price function $p(g_0) = G/g_0$ has negative and increasing first forward differences equal to $-G/\left(g_0(g_0 + 1)\right)$.

We solve numerically under the NBB **Constraint 1** in the case of a uniform distribution with support $[0, qv] = [0, q]$. That is, we set $F(c) = c/(qv) = c/q$.

Fig. 2.11a displays $S_0(g_0, p(g_0))$ for $g_0 = 1, 2, \ldots, 50$ and $G = 13.5, 15$; $\tau=60$, $\lambda=7.5$ and $q = 0.2 < 0.27 = 13.5/50$ so **Constraint 1** holds for all the g_0. The bar plot in dark blue in panel (a) corresponds to the higher value $G = 15$. Only values $g_0 \geq 15$ are then feasible. In any case, the success rate is essentially zero outside the range of g_0 values between 22 and 40. Reducing the price and raising the bidder threshold has an initially positive effect and this is decreasing so that $S_0(g_0, p(g_0))$ is in general concave with a unique maximizer $g_0^* > \underline{g}_0 \triangleq \lceil G \rceil$. For instance, increasing g_0 from 25 to 26 reduces the implied price from 0.6 to 0.577, and that more than compensates the lower success probability caused by a unit increase in the bidder threshold, whereas raising $g_0=30$ to 31 reduces the price only by 0.016 and the success rate drops. If G were sufficiently low (not plotted), p would react too little to g_0 and $g_0^* = \lceil G \rceil$ would be optimal.

---

[12]Note that negative cost bidders would still inspect if $p > v$ since they can inspect and then not bid; with $p > v$, **C** would cease to be optimal but the entrepreneur only values bids and would therefore never set such a price.

(a) $S_0(g_0, p(g_0))$ against g_0 (b) g_0^* against G (c) g_0^* against λ
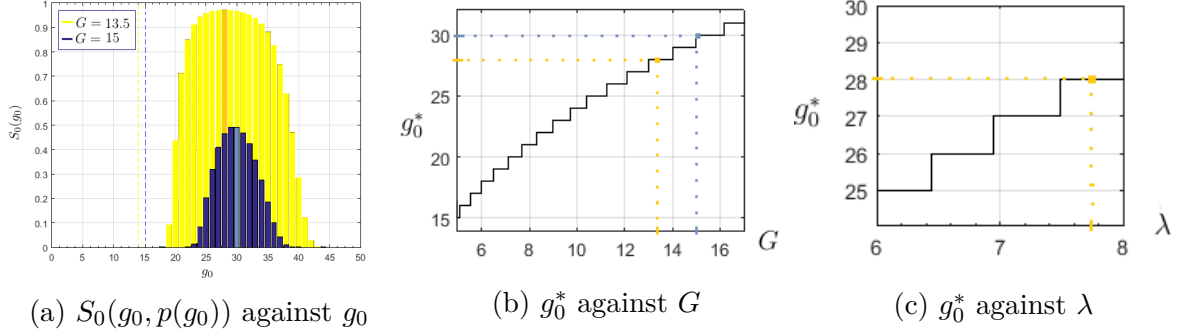
Figure 2.11: Optimal bidder threshold for a campaign with $v = 1$ and inspection costs uniformly distributed on $[0, q]$. Panels (a) and (b) shift G holding the other parameters fixed at $(\tau, \lambda, q) = (60, 7.75, 0.2)$. Panel (c) moves λ and sets $G = 13.5$ and the other parameters as before.

Notice that g_0^* increases from 28 in the case with budget need $G=13.5$ to $g_0^* = 30$ when budget need rises to $G=15$. Fig. 2.11b illustrates the more general property that g_0^* is increasing in G.

Observation 2.1. *The optimal threshold is weakly increasing in financial need G.*

Intuitively, greater financial need and hence a greater funding goal amplifies the price reductions from raising g_0, so raising g_0 is more attractive. Clearly, $g_0^* = 0$ when $G \leq 0$.

To justify the assumption that **Constraint 1** should be satisfied, one possibility is to compare the solution of the restricted problem that satisfies it against an upper bound to the success rate from the solution in the case where it is not satisfied. Instead of completing this verification, we note a simpler solution to implement in the final iteration of this article: if the entrepreneur has a marginal cost per unit produced equal to κ instead of 0 as in the baseline, this immediately obliges her to raise p strictly above κ since her budget constraint becomes $g_0(p - \kappa) \geq G$. Setting $\kappa \geq q$ ensures that **Constraint 1** has to hold. We expect to observe similar results in this case.

2.5.2   Gap-dependent pricing

In this subsection, we fix g_0 so that the only choice variables are the prices charged on each reward tier. The bidding intensity on unit $g_0 - g$ reacts directly to its specific price p_g and indirectly to the price charged on the remaining unit to completion $\mathbf{p}_{g-1}$. For this reason, reducing the price faced by a bidder's successors creates an (indirect) *encouragement effect* that makes the bidder more prone to inspect and bid. In consequence, optimal prices decrease as the gap reduces.

Formally, generalizing problem Eq. (2.39), the objective of the entrepreneur is now to set g_0 and a menu of prices $\mathbf{p}_{g0} = (p_{g_0}, p_{g_0-1}, \ldots, p_1)$ to maximize her success rate subject

to the constraint of reaching her funding goal:

$$\max_{g_0, \mathbf{p}_{g_0}} S_0(\mathbf{p}_{g_0}) \qquad \text{subject to} \qquad G = \sum_{g=1}^{g_0} p_g = v g_0 - \sum_{g=1}^{g_0} d_g \qquad (2.39)$$

The success rate $S_0(\mathbf{p}_{g_0}) \equiv S_{(0,g_0)}(\mathbf{p}_{g_0})$ is defined as the solution of a recursive equation that generalizes Proposition 2.1. Only current and future prices matter so we also define $\mathbf{p}_g \triangleq (p_g, p_{g-1}, \ldots, p_1)$. Again, we focus on the case where NBB holds:

Constraint 2 (Multiprice NBB). *$p_g > qv$ for all g.*

As before, this is guaranteed to hold if the unit cost of production $\kappa \geq q$. Given **Constraint 2**, **C** and **A** are the only relevant substrategies and we again only need to characterize when bidders choose **C**:

Proposition 2.9. *With gap-dependent pricing, the crowdfunding game has a unique PBE $a_{(t,g,c)} = \boldsymbol{C}$ if and only if $c \leq q(v - p_g)S_{(t,g-1)}(\boldsymbol{p}_{g-1})$, giving bid intensity $\beta_{(t,g)}(\boldsymbol{p}_g)$, where*

$$\beta_{(t,g)}(\boldsymbol{p}_g) = \lambda q F\left(q(v - p_g)S_{(t,g-1)}(\boldsymbol{p}_{g-1})\right) \qquad \text{for any } \boldsymbol{p}_g, \; S_{(t,g)}(\boldsymbol{p}_g) = 1 \text{ if } g \leq 0$$

$$(2.40)$$

$$S_{(t,g)}(\boldsymbol{p}_g) = \int_t^\tau \exp\left(-\int_t^{T_{g_0-g+1}} \beta_{(x,g)}(\boldsymbol{p}_g)\, \mathrm{d}x\right) \beta_{(T,g)}(\boldsymbol{p}_g) S_{(T,g-1)}(\boldsymbol{p}_{g-1})\, \mathrm{d}T_{g_0-g+1} \text{ for } g \geq 1$$

$$(\text{REC-}S\text{-}\mathbf{p})$$

Here the stopping time T_n for $n \equiv g_0 - g + 1$ is the date at which the n'th bid is pledged. In the above expressions, the bidding intensity $\beta_{(t,g)}(\mathbf{p}_g)$ depends directly on the specific price p_g that applies to unit $g_0 - g$ and indirectly on the price vector applied to the successive unit sold $\mathbf{p}_{g-1}$ faced by successor bidders. Later we will show that this dependency produces important effects that determine optimal pricing.

Since we fix g_0, the optimization problem is slightly simpler:

$$\max_{\mathbf{p}_{g_0}} S_0(\mathbf{p}_{g_0}) \qquad \text{subject to:} \qquad G = \sum_{g=1}^{g_0} p_g, \quad g_0 \text{ given} \qquad (2.41)$$

We focus on two special cases. A single price design $\mathbf{p}_{g_0}$ is a constant vector of g_0 elements so, in its context, we define $S_0(p, g_0) \equiv S_0(\mathbf{p}_{g_0})$ for brevity. The two-price-tiers scheme discussed previously is instead constituted by g_{0a} units sold initially at price p_a and g_{0b} units sold at prices p_b afterwards, so that total sells to reach the funding threshold are $g_0 \equiv g_{0a} + g_{0b}$. It will be employed for all the numerical exercises that follow. Using linear CDF, This price scheme is fully specified by g_0, g_{0a} and p_a, with p_b following uniquely from the goal constraint $G = p_a g_{0a} + p_b(g_0 - p_a)$. The recursion that determines the success rate, in this case denoted $S_{(t,g)}(g_a, g_b; p_a, p_b)$, follows readily from Eq. (REC-S-$\mathbf{p}$)

(a) Uniform c
$(S_0(15.5/30) = 0.15)$

(b) Homogeneous c
$(S_0(9.5/20) = 0.55)$
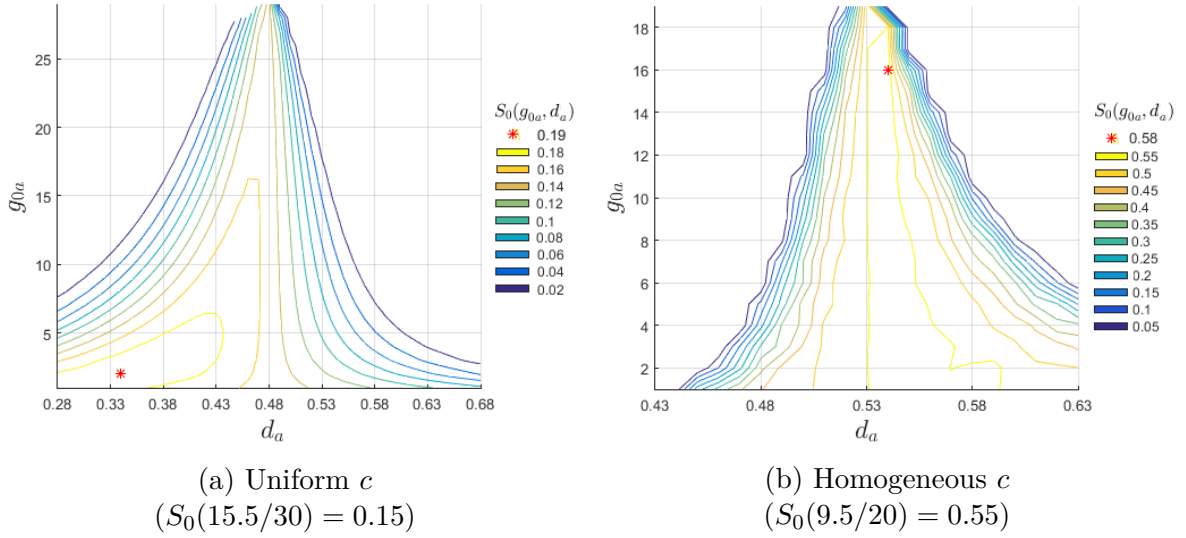
Figure 2.12: $S_0(g_{0a}, d_a)$ isoquants. Panel (a) uses parameters the same parameters as in Fig. 2.11 except for $G = 15.5$. Panel (b) uses parameters $G = 9.5$, $g_0 = 19$, $c = 0.08$, $(\tau, \lambda, q) = (60, 3, 0.2)$.*

*Isolines only connect points.

after adjusting the bid intensity as

$$
\beta_{(t,g)} = \begin{cases}
\lambda q F\left(q(v - p_a)S_{(t,g-1)}\left(g_{0a} - (g_0 - g + 1), g_{0b}; p_a, p_b\right)\right) & \text{for } g > g_0 - g_{0a} \\
\lambda q F\left(q(v - p_a)S_{(t,g_0-g_{0a}-1)}\left(0, g; p_a, p_b\right)\right) & \text{for } g = g_0 - g_{0a} \\
\lambda q F\left(q(v - p_b)S_{(t,g-1)}\left(0, g - 1; p_a, p_b\right)\right) & \text{for } g < g_0 - g_{0a}
\end{cases}
$$

We determine numerically the optimal two-tier pricing using the same uniform CDF as in the single price design and find that $p_a^* > p_b^*$. Basically, lower prices in phase b are optimal because they encourage bidders arriving in phase a.

Observation 2.2. *Assuming uniform CDF, price frontloading is optimal: the QLR price scheme involves $p_a^* > p_b^*$*

Fig. 2.12a presents our numeric result showing the success isoquants and the optimal menu in the (d_a, g_{0a}) coordinates system with G=15.5 and g_0=30. Forcing initial higher prices raises the success rate of $S_0(g_0, p)$=0.15 achieved with a single price $p = 0,52$ to $S_0(g_{0a}, g_{0b}; p_a, p_b)$=0.19 by selling $g_{0a} = 2$ units at $p_a = 0,66$ and $g_{0b} = 28$ units at $p_b = 0,49$. Conversely, price schemes with initial discounts reduce the probability of success, so if initial higher prices are forbidden, the entrepreneur sets a single price.

The optimality of higher initial prices under uniformly distributed inspection costs can be easily proved in a simplified setting where $n = g_0$ bidders arrive according to a deterministic sequence. The same logic applies to continuous time.

Lemma 2.8. *Suppose $n = g_0$, bidders arrive sequentially and inspection costs are uniformly distributed over $[0, q]$. Then optimal prices are frontloaded (increasing in the gap g):*

$$p_g^* = v - \frac{2(vg_0 - G)(g_0 - g + 1)}{g_0(g_0 + 1)}, \quad g_0 \geq g \geq 1 \tag{2.42}$$

The **Proof** *in Appendix 2.A* determines the optimal prices solving the Lagrangian associated to the success rate maximization problem (2.41).

Optimal prices reduce linearly in the number of prior bidders due to the encouragement effect. Taking into account that units are charged with $p_g = v$ after completion, a gap decreasing price scheme for the first g_0 units sold implies an overall non-monotonic price scheme. This non-monotonicity produces a drop in sales after the goal is reached caused by higher prices as observed in crowdfunding data.

The encouragement effect however not always dominates the substitution effect of having to raise initial prices. Whether it dominates depends on the price elasticity of inspection incentives. When inspection costs are homogeneous complementary originates from shifts in the critical dates $\hat{t}_g(\mathbf{p}_{g-1})$, hence the encouragement effect is weaker than with a uniform CDF, so the substitution effect prevails more easily. In some cases the encouragement is shut-down completely, e.g. when $n = g_0$ bidders arrive sequentially and all of them are needed to reach success. Setting $\lambda = v = 1$, we have $S_0 = q^n$ and $S_{(t,g)}^{\text{bid}}$ along feasible paths is either 0 or q^{g-1}, so the optimal price is $p_g^* = 1 - c/(q^g)$ such that $S_{(t,g)}^{\text{bid}} d_g q = c$. Higher prices and inspection of predecessor bidders are clear substitutes. When substitution effects prevail, optimal prices decrease with the gap, so we obtain the diametrically opposite result than with the uniform CDF.

Observation 2.3. *Initial discounts are sometimes optimal under homogeneity.*

Fig. 2.12b shows success isoquants assuming homogeneous inspection cost and for $G = 9.5$ and $g_0 = 19$. The success rate increases from $S_0(g_0, p) = 0.55$ with a fixed price $p = 0,47$ to $S_0(g_{0a}, g_{0b}; p_a, p_b) = 0.58$ by selling the first $g_{0a} = 16$ units at price $p_a = 0,46$ and the remaining $g_{0b} = 3$ units at price $0,51$. Moreover, all price menus obtained by setting (g_{0a}, d_a) combinations within the highest yellow isoquant improve the fixed price design.

2.6 Discussion

This section adds further discussion on modeling choices, testable implications of our theory and crowdfunding design results.

2.6.1   Model remarks

Alternative assumptions: We discuss assumptions on the campaign duration τ and bidder inspection cost CDF F. The duration τ is in principle a choice variable although the vast majority of Kickstarter campaigns set it at 30 days (the second most popular option is to set $\tau = 60$ days). If chosen independently within our model with a fixed arrival rate and bidder type distribution, the optimal campaign duration is infinite due to its obvious effect on success rates, profits and consumer rents. It is easy to introduce a trade-off with bidder incentives (and the entrepreneur's profits) by adding generic time discounting. The are also plenty of other factors that justify a trade-off between duration τ and arrival rate λ; e.g. product-specific discounting from boredom, risks of changing tastes and opportunity costs, the fall in arrival rates from increased inter-project competition and increased costs of any ongoing advertising.

Despite restricting bids to 0 or p, the analysis applies for campaigns that accept donations, as standard on Kickstarter, if the restriction does not bind. It also applies if we relax the assumption of unit demand funders if they are purely motivated by consumption and price discounts do not exceed 50%; i.e. if $v < 2p$. It might seem that this is true even without bid restrictions, because donating a fraction of p does not reduce the number of bids that other bidders need to provide given the funding goal is set at an integer multiple of p. However, there could exist equilibria in which multiple bidders donate say $1.5p$ in the hope that there will be multiple such donations, which combine to reduce the gap.

Finally, when we assume F is polynomial we do not allow for other lower order terms than the degree of F apart from the 0-type atom $z^0 \equiv z$, e.g. a linear interaction when discussing $\rho = 2$. Full parsimony would require them though their interpretation is less clear.

Alternative normalizations: Given that τ and λ have equivalent effects on bid profiles and success rates, a natural normalization is to interchange t with relative time $t/\tau \in [0, 1]$ and scale λ to $\lambda\tau$. Even setting $\tau = 1$ is plausible. We did not opt for applying this normalization since remaining time expressed as $1 - t$ instead of $\tau - t$ makes expressions more obscure. These normalizations are clearly not valid if we assume endogenous attention. In the first circulated version of this paper, we recorded time in terms of $r = \tau - t$ but time moving forward is more familiar for describing graphs and linking slopes to derivatives.

Deb et al. (2021); Alaei et al. (2016) assume only a single source of stochasticity in bidding behaviour; either arrivals or valuations while keeping inspection costs homogeneous. Two ways to reduce stochasticity in our setting are to set $q = 1$ as explained above or use the neutral arrival rate $\lambda = 1$. These two normalizations are not analogous: q enters in the inspection incentive directly and indirectly through S while λ enters only indirectly.

2.6.2 Testable implications

Our theory is broadly consistent with stylized facts on the dynamics of crowdfunding purchases even though we abstract from donations. The theory is tractable and we discuss how extensions can readily generate testable predictions. In particular, we predict a positive correlation between duration and funding goals, as observed in the data.

Our main predictions are a decreasing bid profile from decreasing pivotality and a potential increasing tendency when the upside of bidding in success uncertainty dominates the downside of negative news. Our findings provide a novel set of testable implications. The analysis of DPE and JE reveals that the bid profile is differentiated across project categories with different features. The analysis of DPE shows that the initial decreasing slope is steeper for large g_0/τ. This can be tested with a breakdown by campaign category of pledge data. Our theory predicts that campaign with big funding goals, e.g. technology, design, games and film&video, have a steeper initial negative slope due to the stronger DPE than projects with smaller goals, such as for the music, comics, dance and crafts categories. We also predict profile features based on completion time and success rate by category. The bid profile predicted by our model for early completers is near to flat. These campaigns encompass those having strong impact on future success prospects of the entrepreneur so that she sets a low funding goal to ensure a higher success rate. On the other hand, we predict a more pronounced U-shape for successful late finisher campaigns, and a strong decreasing pattern with a slight final increase for unsuccessful campaigns. We expect bid profiles to slope downwards when inspecting is an attractive distraction for most relevant bidders. Conversely, positive slopes are likely when inspection costs are too costly for most bidders to want to inspect. We also predict that projects that spark bidders' curiosity, active enjoyment in reading or desire for distraction tend to outweigh their costs of effort and opportunity cost of time. In this case, negative net inspection costs predominate. Also close contacts and family of the entrepreneur are likely to be anyway informed or to feel obliged to pay attention. Bidder connectedness with the entrepreneur, as estimated via Facebook links or geography (see Agrawal et al., 2011), can serve as a proxy for negative inspection costs.

From the quantitative point of view, one weakness of our theory is that interim slopes are not sharp enough to explain the "peaks" in revenues at the campaign extremes. The final peak in purchases observed in the data is likely caused by bidders who strategically time their inspection decision so as to face less success rate uncertainty. By waiting, they reduce their uncertainty though they oblige others to face more uncertainty or a lower success rate. In a companion paper, we develop a thorough model of the dynamic free-riding that results. Wndogenous timing readily generates a spike of bidding at the last moment, just before the campaign deadline. Strategic waiting by bidders is not the only explanation for the final bunching: also donors make last minute contribution

to ensure the campaign succeeds if still far from reaching the goal as shown by Deb et al. (2021) and Crosetto and Regner (2018). Behavioural theories attribute bunching at the end of campaigns to procrastination and deadline effects. Basically, irrationality is caused by laziness or anxiety under the pressure of time running out. Another possible cause is endogenous advertising intensity from both the crowdfunding platform and the entrepreneur.

An empirical puzzle illustrated by Deb et al. (2021) that merits our demonstrating is that bidding drops at success time (T_{g_0}) and tails off after (even for purchases). In section Section 2.5, this occurs because prices increase. Without endogenous pricing, relative bidding prior and post success depends on when success occurs. If we only look at successes, then we have g_0 bids by T_{g_0}, so average bidding rate pre-success equals g_0/T_{g_0}. Average bidding rate after success is instead λq (given we only look at successes here). So if τ is small: $g_0/\tau > \lambda q$, so average bidding definitely falls after success. If τ is large, then pre- and post- success bidding comparison is just dependent on T_{g_0} in the obvious way. The expectation over the g_0 distribution can result in either possibility. If we also sample unsuccessful campaigns, so that the average bidding before T_{g_0} coincides with average bidding of failed campaigns, then average bidding is likely to be higher after success. The reason is that pre-success average puts more weight on failed campaign (with lower bidding), while post- only includes successes.[13]

2.6.3   Crowdfunding design

Typical crowdfunding campaigns feature a diversity of rewards that offer variants on the entrepreneur's product at differing prices. In addition, the quantity of some of these rewards may be limited. We refer to a crowdfunding reward design with these features as using a *quantity limited reward* (QLR) scheme. Sometimes quantity limits are driven by non-linear production costs, as when, for example, a music band that has only time to provide a limited number of private concerts in return for high price bidders. Quantity limits may also aim to create exclusivity. This practice is pervasive so it is its optimal design deserves theoretical understanding. However, Kickstarter prohibits modifications of the reward tiers during the campaign (after the first unit of a tier is sold) making QLR static price menus. Hence, adapting our dynamic framework to encompass QLR requires adding an additional layer of complexity. We now discuss how our results on dynamic pricing via gap-dependency extend to having bidders choosing prices from a simultaneous menu.

QLR with initial discounts is a common practice in reality and constitutes effectively a form of gap-dependent pricing. This scheme involves a form of early-bird discounts in

---

[13]For the homogeneous cost case, average bidding falls after success even conditional on success, but for this reason the slope in Fig. 2.9a for the binary cost case goes down more when controlling for success.

that gaps are higher for earlier arrivals in a relative sense. The entrepreneur can therefore implement a gap decreasing price due to the first-come-first-served logic without directly obliging the desired gap-dependent pricing.

Also the converse case where bidders buy a product at the highest available price is possible. Ellman and Hurkens (2019b) note that there are essentially two ways to design a crowdfunding campaign that offers the same product at different prices. The first is to charge higher prices and include an added good that is sometimes only symbolic, such as a signature. The second, is to associate to the basic reward a donation that constitutes an overprice, optimally with a minimum amount specified. A static price menu with these features corresponds to gap dependent pricing if bidders willingly pay the initial higher prices. This requires low initial success prospects relative to their valuation of the crowdfunding product.

In both discussed cases, we have to take into account that the gap threshold g_0 might not apply. Also, if prices fall as the gap falls, the commitment to reduce prices later is not easy even if the entrepreneur is able to add apparently new rewards. E.g. suppose $G = 5$ and the price menu has two price tiers, a and b with $p_b = 2$ and $p_a = 1$. If bidders the first two bidders buy units from tier a, then the lowest bidder threshold that achieves the goal is 4 whereas it is 3 if they first buy from tier b. Due to this complication, a comprehensive treatment of menu price schemes requires a fundamental change in the model mechanics that is outside the scope of this article.

After-market: In our baseline model of Section 2.2, we assumed bidding was now or never. There is no alternative to buy the product in an *after-market* following a successful conclusion to a crowdfunding campaign. With a single positive valuation, the option to wait until then is neutral to crowdfunding dynamics. Assuming inspection costs drop to zero when the campaign is over, the entrepreneur would then extract the bidder rent fully setting $p = v$ given our assumptions. So bidders would be indifferent between waiting and playing A. Of course, that is a very special case but the more general point is that crowdfunding creates a price discontinuity. With multiple positive valuations, the option of after-market purchases stops being neutral because bidders can now hope to gain a rent from buying in the after-market. The rent was always zero in the unique positive valuation case because then the entrepreneur would always set price equal to that valuation.

Advertising: Throughout, we fixed parameters that depend on advertising, promotion and selection strategies of platforms and entrepreneurs. Endogenizing those parameters (bidder arrival intensity λ, product quality variables v, q and campaign duration τ) lies beyond the scope of this paper since it requires a model of the advertising technology and attention competition between projects. The types of tradeoff derived in Section 2.5 will continue to apply.

2.7   Conclusion

The main contribution of this paper is identification of the decreasing pivotality and Jensen effects which will apply in a much broader range of dynamic participation games than just crowdfunding. At the same time, we have presented a parsimonious model of crowdfunding that is able to account for a variety of observed momentum effects in the dynamics of funding, including the much-commented U-shape profile.

Our baseline model focuses on pure private values and bidders exogenously constrained to either forget the project or inspect and possibly bid on it in the same decision episode in which they become aware of the project. Reality is much more complex, but the insights from our basic set-up readily extend to cover richer scenarios and to generate additional effects such as free-riding in the endogenous moves setting.

The simplicity of the basic model served to pinpoint two clear driving forces, despite the relatively complex dynamic challenge faced by bidders. They must anticipate how project success chances depend on the inspection calculus of all bidders who have not yet moved. Leaving aside friends, fans and random procrastinators who enjoy looking at the entrepreneur's project, each bidder is only willing to dedicate time or effort to inspecting a project if learning his type is likely to be instrumentally valuable. That, in turn, requires the project to have a reasonable chance of success. As a result, the distribution of inspection costs and bidder knowledge about those costs take centre-stage.

Despite its parsimony, the model is broadly consistent with the U-shaped dynamics that initially motivated our study and with the stylized facts presented in previous literature. The model can account for the most salient facts, generating a U-shaped profile, as well as increasing and decreasing patterns, depending on the project's starting condition and the distribution of inspection costs. Nonetheless, alternative and complementary theories certainly could provide plausible explanations for some aspects of the U-shape via assumptions about the advertising technology and word-of-mouth communication. Advertising and platform promotion strategies represent an important part of the puzzle, but to derive the optimal promotion strategies, we first want to know how the dynamics of campaigns respond to an exogenous bidder arrival rate. We have shown how the threshold alone can interact with costly inspection to generate a range of dynamic effects that will continue to be relevant in richer models. Other forces, such as word-of-mouth learning and common value effects, can influence the observed funding patterns and are likely to be empirically important as well.

Advertising and project promotion effects are highly compatible with our model. An alternative story for the profile is that initial advertising leads to an increasing profile when amplified via word-of-mouth (before saturation sets in). Towards the end of the campaign, the opposite may occur as advertising influence decays over time. Depending on the specific technology assumption, it is likely that a lot could be gained from

strategically-timed promotions that raise bidder arrival rates precisely when projects are at risk of going cold and dying out. Such an analysis would have to pay attention to bidders' awareness of the promotion strategies in use. Competing campaigns also merit investigation.

The role of common values also merits further investigation. Common values readily generate positive and negative cascades but combining a non-extreme common values model with the threshold effect in our dynamic framework is far from trivial.

Our design results make the analysis relevant for entrepreneurs who need to evaluate how their strategic choices affect project success prospects. They know that raising price lowers the chance of a success. Our analysis shows that the resulting trade-off is significantly more complicated than simply assessing the initial gap as sufficed in the static context with a single price. Dividing the gap by time available is a natural first look in a setting where bidders arrive over the duration of the project, but inspection costs can make many potential bidders useless: they just pass over the project if the gap is too high relative to time left on the campaign clock.

Appendices of Chapter 2

2.A   Proofs

Proof of Lemma 2.1. Bid intensity $\beta_{(t,g)}$ cannot be negative so (i) and Eqs. (2.2) and (ODE-S) immediately imply (ii). So we need only prove $\Delta S_{(t,g)} \geq 0, \forall g$; this implies the second inequality in (i) since $S_{(t,g)}^{\text{bid}} \equiv S_{(t,g-1)}$. We prove $\Delta S_{(t,g)} \geq 0$ by induction. First, if $g \leq 0$, $S_{(t,g)} \equiv 1$ so $\Delta S_{(t,g)} = 0$. Second, $\Delta S_{(t,1)} = 1 - S_{(t,1)} = e^{-\lambda q(\tau - t)} \in (0,1]$. For $g \geq 2$, we integrate Eq. (2.8) by parts, using $\mathrm{d}n_T = -n_T \beta_{(T,g)} \, \mathrm{d}T$ from Eq. (2.7), to give

$$S_{(t,g)} = -\int_t^\tau S_{(T,g-1)} \, \mathrm{d}n_T = \int_t^\tau n_T \dot{S}_{(T,g-1)} \, \mathrm{d}T - \left[n_\tau S_{(\tau,g-1)} - n_t S_{(t,g-1)} \right] \qquad (2.A.1)$$

$S_{(\tau,g-1)} = 0 \, \forall g \geq 2$ and $n_t = e^0 = 1$ imply $\left[n_T S_{(T,g-1)} \right]_t^\tau = S_{(t,g-1)}$. So using Eq. (ODE-$S$),

$$\Delta S_{(t,g)} = -\int_t^\tau n_T \dot{S}_{(T,g-1)} \, \mathrm{d}T = \int_t^\tau n_T \beta_{(T,g-1)} \Delta S_{(T,g-1)} \, \mathrm{d}T$$

Now $\beta_{(t,g)} \geq 0$ and $n_x > 0 \, \forall x \leq \tau$, so (i) holds at $g - 1 \implies$ (i) holds at g. ∎

Proof of Lemma 2.3 (ii).

$$
\dot{A}_t^{\mathcal{S}} = \frac{1}{S_0} \sum_{g=-\infty}^{K} \left(\dot{Q}_{(t,g)} \beta_{(t,g)} S_{(t,g-1)} + Q_{(t,g)} \dot{\beta}_{(t,g)} S_{(t,g-1)} + Q_{(t,g)} \beta_{(t,g)} \dot{S}_{(t,g-1)} \right)
$$

$$
= \frac{1}{S_0} \sum_{g=-\infty}^{K} \left(\left(Q_{(t,g+1)} \beta_{(t,g+1)} - Q_{(t,g)} \beta_{(t,g)} \right) \beta_{(t,g)} S_{(t,g-1)} + Q_{(t,g)} \dot{\beta}_{(t,g)} S_{(t,g-1)} + Q_{(t,g)} \beta_{(t,g)} \dot{S}_{(t,g-1)} \right)
$$

$$
= \frac{1}{S_0} \sum_{g=-\infty}^{K} Q_{(t,g)} \left(\beta_{(t,g)} \left(\beta_{(t,g-1)} S_{(t,g-2)} - \beta_{(t,g)} S_{(t,g-1)} \right) + \dot{\beta}_{(t,g)} S_{(t,g-1)} + \beta_{(t,g)} \dot{S}_{(t,g-1)} \right)
$$

by change of variable from $g+1$ to g in the first summation term. Next from (ODE-S),

$$
\dot{A}_t^{\mathcal{S}} = \frac{1}{S_0} \sum_{g=-\infty}^{K} Q_{(t,g)} \left(\beta_{(t,g)} \left(\beta_{(t,g-1)} - \beta_{(t,g)} \right) + \dot{\beta}_{(t,g)} \right) S_{(t,g-1)}
$$

$$
= \frac{1}{S_0} \sum_{g=-\infty}^{K} Q_{(t,g)} \mathcal{L}_{(t,g)}^{\beta} S_{(t,g-1)}
$$

For $\dot{A}_t^{\mathcal{F}}$, note that $A_t^{\mathcal{F}} = \frac{1}{1-S_0} \left(A_t - S_0 A_t^{\mathcal{S}} \right)$, so

$$
\dot{A}_t^{\mathcal{F}} = \frac{1}{1-S_0} \left(\dot{A}_t - S_0 \dot{A}_t^{\mathcal{S}} \right)
$$

$$
= \frac{1}{1-S_0} \sum_{g=-\infty}^{K} Q_{(t,g)} \mathcal{L}_{(t,g)}^{\beta} \left(1 - S_{(t,g-1)} \right)
$$

∎

Proof of Lemma 2.4. We have F' with lower costs than F in that $F \succeq_{\text{FOSD}} F'$. Let H', H denote the CDF of T_g under F', F, respectively and $\mathbb{E}(S_{(T,g-1)}|H)$ indicates expectation over $T_g \equiv T$ distributed according to H. Similarly, we use S', β' and S, β to distinguish results for F' and F. By the inductive hypothesis at $g-1$, $S'_{(t,g-1)} \geq S_{(t,g-1)}$ so $\beta'_{(t,g)} \geq \beta_{(t,g)}$ by Eq. (2.9) and so $n'_T = \exp\left(-\int_t^T \beta'_{(x,g)}(z)\, \mathrm{d}x \right) \leq n_T$ for all T. For any t, g (suppressed), $H_T \equiv 1 - n_T$ so $H' \geq H$ for all T; i.e., $H \succeq_{\text{FOSD}} H'$.

Now, S, S' are decreasing in T by Lemma 2.1 so by FOSD, $\mathbb{E}(S_{(T,g-1)}|H') \geq \mathbb{E}(S_{(T,g-1)}|H)$; to prove this FOSD result, we integrate by parts as in Eq. (2.A.1) ($\dot{S} \leq 0$ by Lemma 2.1):

$$
\int_t^\tau S_{(T,g-1)}\, \mathrm{d}H_T' - \int_t^\tau S_{(T,g-1)}\, \mathrm{d}H_T = S_{(\tau,g-1)} \left(H_\tau' - H_\tau \right) - \int_t^\tau \dot{S}_{(T,g-1)} (H_T' - H_T)\, \mathrm{d}T
$$

since $H_t^{(t,g)} = 0 \implies H_t = H_t' = 0$ and $S_{(\tau,g-1)} = S'_{(\tau,g-1)}$ ($= 0$ if $g \geq 2$, $= 1$ if $g = 1$). Finally, applying Eq. (REC-S) and then the inductive hypothesis,

$$
S'_{(t,g)} - S_{(t,g)} = \mathbb{E}(S'_{(T,g-1)}|H') - \mathbb{E}(S_{(T,g-1)}|H) \geq \mathbb{E}(S_{(T,g-1)}|H') - \mathbb{E}(S_{(T,g-1)}|H) \geq 0
$$

∎

Proof of Proposition 2.4. The chain rule for generators (see Appendix 2.C.1) states that

$$\mathcal{L}^{h(Y)}_{(t,g)} = h_Y \mathcal{L}^{Y}_{(t,g)} + \beta_{(t,g)} \left(\Delta h(Y_{(t,g)}) - h_Y \Delta Y_{(t,g)} \right) \tag{2.A.2}$$

Setting $Y_{(t,g)} \equiv q S^{\text{bid}}_{(t,g)}$ and $h(\cdot) \equiv \lambda q F(\cdot)$ so that $\beta_{(t,g_t)} \equiv h(Y_{(t,g_t)})$ and $h_Y(\cdot) \equiv \lambda q F_c(\cdot)$, provides an expression for $\mathcal{L}^{\beta}_{(t,g)}$ which we equate with its DPE-JE decomposition (2.19):

$$\lambda q \left(q F_c \, \mathcal{L}^{S^{\text{bid}}}_{(t,g)} + \beta_{(t,g)} \left(\Delta F(q S^{\text{bid}}_{(t,g)}) - F_c \, q \Delta S^{\text{bid}}_{(t,g)} \right) \right) = \lambda q \left(\mathcal{E}^{(t,g)} + \mathcal{J}^{(t,g)} \right) \tag{2.A.3}$$

where $F_c \equiv F_c(q S^{\text{bid}}_{(t,g)})$. Cancelling $\mathcal{E}^{(t,g)} = q F_c \, \mathcal{L}^{S^{\text{bid}}}_{(t,g)}$ by Eqs. (2.13) and (2.15) proves Eq. (2.17):

$$\mathcal{J}^{(t,g)} = \beta_{(t,g)} \left(\Delta F(q S^{\text{bid}}_{(t,g)}) - F_c \, q \Delta S^{\text{bid}}_{(t,g)} \right)$$

If F is polynomial with maximal exponent $\rho \in \mathbb{N}_+$,

$$\Delta F(q S^{\text{bid}}_{(t,g)}) = F_c(q S^{\text{bid}}_{(t,g)}) q \Delta S^{\text{bid}}_{(t,g)} + \frac{F_{cc}(q S^{\text{bid}}_{(t,g)})}{2} \left[q \Delta S^{\text{bid}}_{(t,g)} \right]^2 + \sum_{k=3}^{\rho} \frac{D^k_c(F) \left(q S^{\text{bid}}_{(t,g)} \right) \left[q \Delta S^{\text{bid}}_{(t,g)} \right]^k}{k!}$$

The instantaneous variance of $S^{\text{bid}}_{(t,g)}$ equals the intensity of the underlying Poisson process g_t times the jump size squared, so $\nu^{q S^{\text{bid}}}_{(t,g)} = \beta_{(t,g)} (q \Delta S^{\text{bid}}_{(t,g)})^2 = q^2 \nu^{S^{\text{bid}}}_{(t,g)}$ (details in 2.C.4). If $\rho = 2$, this implies $\mathcal{J}^{(t,g)} = \frac{q^2}{2} \nu^{S^{\text{bid}}}_{(t,g)} F_{cc}(q S^{\text{bid}}_{(t,g)})$. So Eq. (2.19) implies that $\mathcal{L}^{\beta}_{(t,g)}$ satisfies (2.20) and condition (2.21) guarantees $\mathcal{J}^{(t,g)} \geq \mathcal{E}^{(t,g)}$, hence $\mathcal{L}^{\beta}_{(t,g)} \geq 0$. If $\rho \geq 3$,

$$\mathcal{J}^{(t,g)} = \frac{q^2}{2} \nu^{S^{\text{bid}}}_{(t,g)} F_{cc}(q S^{\text{bid}}_{(t,g)}) + \beta_{(t,g)} \sum_{k=3}^{\rho} \frac{D^k_c(F(c)) \left[q \Delta S^{\text{bid}}_{(t,g)} \right]^k}{k!} \tag{2.A.4}$$

k'th derivatives for $k > \rho$ are zero. If $F(c) \equiv \sum_{k' \in \{0,1,...,\rho\}} c_{k'} c^{k'}$ for $3 \leq k \leq k' \leq \rho$, the k'th summation term from positive polynomial coefficient $c_{k'}$ is a positive multiple of $c^{k'-k}$. So condition (2.21) is sufficient for the JE to prevail over the DPE and produce an increasing expected bid rate from state (t,g); if $\rho = 2$, (2.21) is also necessary. ∎

Proof of Lemma 2.6. Substituting for $\nu^{S^{\text{bid}}}_{(t,g)}$, $\mathcal{L}^{S^{\text{bid}}}_{(t,g)} = \mathcal{L}^{\Delta S}_{(t,g)}$ from equation Eq. (2.13) and the expressions for F_c, F_{cc} from just below Eq. (2.30), Eq. (2.20) reduces to

$$\mathcal{L}^{\beta}_{(t,g)} = \lambda q (1-z) \Delta S^{\text{bid}}_{(t,g)} \left(-2 S^{\text{bid}}_{(t,g)} \Delta \beta_{(t,g)} + \beta_{(t,g)} \Delta S^{\text{bid}}_{(t,g)} \right)$$

Simplifying further using Eq. (2.6) and $(S^{\text{bid}}_{(t,g-1)})^2 - (S^{\text{bid}}_{(t,g)})^2 = \Delta S^{\text{bid}}_{(t,g)} (S^{\text{bid}}_{(t,g-1)} + S^{\text{bid}}_{(t,g)})$,

$$\mathcal{L}^{\beta}_{(t,g)} = (\lambda q)^2 (1-z) \left(\Delta S^{\text{bid}}_{(t,g)}(z) \right)^2 \left[z - (1-z) S^{\text{bid}}_{(t,g)}(z) \left(2 S^{\text{bid}}_{(t,g-1)}(z) + S^{\text{bid}}_{(t,g)}(z) \right) \right]$$

The term in square brackets determines the generator's sign. It is positive under condition

(2.31), strictly so for all $g \geq 2$. Maximizing $S^{\text{bid}}_{(t,g)}, S^{\text{bid}}_{(t,g-1)}$ at 1 gives the highest lower bound on z at $z = 3/4$. ∎

Proof of Lemma 2.8.   The Lagrangian of the maximization problem (2.41) is given by:

$$L = \left(\prod_{g=1}^{g_0} d_g^{g_0-g+1} \right) - l \left[\sum_{g=1}^{g_0} d_g - (vg_0 - G) \right] \tag{2.A.5}$$

where l is the multiplier associated with the goal constraint. The first-order condition with respect to generic d_g is

$$\prod_{g=1}^{g_0} d_g^{g_0-g+1} \left(\frac{g_0 - g + 1}{d_g} \right) = l \quad \text{So,} \quad \frac{d_g}{d_1} = \frac{g_0 - g + 1}{g_0} \tag{2.A.6}$$

Now using Eq. (2.A.6) together with the goal constraint goal we find the exact value of d_1:

$$vg_0 - G \quad = \frac{d_1}{g_0} \left[\sum_{g=1}^{g_0} g_0 - g + 1 \right] \quad = \frac{d_1}{g_0} \left[g_0(g_0 + 1) - \frac{g_0(g_0 + 1)}{2} \right] \quad = \frac{d_1(g_0 + 1)}{2}$$

which implies

$$d_1 = \frac{2\,(vg_0 - G)}{g_0 + 1} \quad \text{and} \quad d_g = \frac{2\,(vg_0 - G)\,(g_0 - g + 1)}{g_0(g_0 + 1)}$$

The optimal prices in Eq. (2.42) follow by subtracting this d_g from v. ∎

2.B   Derivations for canonical distributions

2.B.1   CDF equivalence

Any CDF $F(c)$ and arrival rate λ is outcome-equivalent to the CDF $F'(c)$ on support $[0, q]$ with rescaled arrival rate $\lambda' = \lambda F(q)$ where $F'(c)$ is derived from $F(c)$ by truncating values $c > q$ and substituting values $c < 0$ by 0: $F'(c)$ has atom $z = F(0)/F(q)$ at zero and

$$F'(c) \triangleq z + (1 - z)\frac{F(c) - F(0)}{F(q) - F(0)} \qquad \text{for } c \in [0, q]$$

$(\lambda, F(c))$ and $(\lambda', F'(c))$ generate identical bid profiles.

2.B.2 Average bidding for $g_0 = 2$ under linear CDF

With $g_0 = 2$, $\beta_{(t,1)}$ is constant at λq and $\beta_{(t,2)}$ equals $\lambda q F(q S_{(t,1)})$. Averaging this,

$$A_t = \lambda q \left[1 + Q_{(t,2)} \left(F\left(q S_{(t,1)} \right) - 1 \right) \right] \tag{2.B.1}$$

$$= \lambda q \left[1 + Q_{(t,2)} \left(z + (1 - z) S_{(t,1)} - 1 \right) \right] \tag{by 2.28}$$

$$= \lambda q \left[1 + \exp\left(-\lambda q \int_0^t z + (1 - z) S_{(x,1)} \, \mathrm{d}x \right) \left(z + (1 - z) S_{(t,1)} - 1 \right) \right] \tag{by 2.24}$$

$$= \lambda q \left[1 - \exp\left(-\lambda q \left(zt + (1 - z) \int_0^t S_{(x,1)} \, \mathrm{d}x \right) \right) (1 - z) \left(1 - S_{(t,1)} \right) \right] \tag{2.B.2}$$

Again from (2.24),

$$S_{(t,1)} = 1 - e^{-\lambda q(\tau - t)} \implies \int_0^t S_{(x,1)} \, \mathrm{d}x = t - \frac{1}{\lambda q} \left(e^{-\lambda q(\tau - t)} - e^{-\lambda q\tau} \right) \tag{2.B.3}$$

and finally combining Eqs. (2.B.2) and (2.B.3),

$$A_t = \lambda q \left\{ 1 - (1 - z) \exp\left[-\lambda q\tau + \lambda q(1 - z) \left(e^{-\lambda q(\tau - t)} - e^{-\lambda q\tau} \right) \right] \right\} \tag{2.B.4}$$

Intuitively, parameters $(\lambda q, z, \tau)$ raise bidding, as clear from expression (2.B.4):

2.B.3 Comparative statics on z for $g_0 = 2$ under linear CDF

For any $g_t \leq 1$, bid intensity is constant at λq. Whereas while $g_t = 2$, the rate of change in the bid intensity follows from applying the formula of the rate of decrease in pivotality, from Eq. (2.12), to the generator in Eq. (2.29):

$$\mathcal{L}^{\beta}_{(t,2)} = -\lambda q(1 - z) \, \Delta S_{(t,1)} \, \Delta \beta_{(t,2)}$$

$$= -(\lambda q)^2 (1 - z) \left[1 - \left(z + (1 - z) S_{(t,1)} \right) \right] \left(1 - S_{(t,1)} \right)$$

$$= - \left(\lambda q(1 - z) \right)^2 \left(1 - S_{(t,1)} \right)^2$$

$$= - \left(\lambda q(1 - z) \right)^2 e^{-2\lambda q(\tau - t)}$$

Since $g_0 = 2$, we need the probability $Q_{(t,2)}$ that $g_t = 2$ from Eq. (2.24),

$$Q_{(t,2)} = \exp\left[-\int_0^t \lambda q\left(z + (1-z)S_{(x,1)}\right)\mathrm{d}x\right]$$

$$= \exp\left[-\lambda q\int_0^t z + (1-z)\left(1 - e^{-\lambda q(\tau-x)}\right)\mathrm{d}x\right] = \exp\left[-\lambda q\left(t - (1-z)\int_0^t e^{-\lambda q(\tau-x)}\,\mathrm{d}x\right)\right]$$

$$= \exp\left[-\lambda q\left(t - (1-z)\frac{e^{-\lambda q(\tau-t)} - e^{-\lambda q\tau}}{\lambda q}\right)\right] = \exp\left[-\lambda q t + (1-z)\left(e^{-\lambda q(\tau-t)} - e^{-\lambda q\tau}\right)\right]$$

It follows that, starting from $g_0 = 2$,

$$\dot{A}_t = Q_{(t,2)}\,\mathcal{L}_{(t,2)}^{\beta} = -\left(\lambda q(1-z)\right)^2 \exp\left[-\lambda q(2\tau - t) + (1-z)\left(e^{-\lambda q(\tau-t)} - e^{-\lambda q\tau}\right)\right]$$

So the slope's magnitude is falling in z as shown by the partial derivative,

$$D_z|\dot{A}_t| = -(\lambda q)^2(1-z)\exp\left[-\lambda q(2\tau - t) + (1-z)\left(e^{-\lambda q(\tau-t)} - e^{-\lambda q\tau}\right)\right] \qquad (2.B.5)$$

$$\times\ \left(2 + (1-z)\left(e^{-\lambda q(\tau-t)} - e^{-\lambda q\tau}\right)\right) \le 0$$

2.B.4   Condition for upward slope for quadratic CDF and $g_0 = 2$

For $g_0 = 2$, similar to the affine case, the slope is determined by the product of

$$\mathcal{L}_{(t,2)}^{\beta} = (\lambda q)^2(1-z)e^{-2\lambda q(\tau-t)}\left[z - (1-z)\left(1 - e^{-\lambda q(\tau-t)}\right)\left(3 + e^{-\lambda q(\tau-t)}\right)\right] \quad (2.B.6)$$

$$\text{and } Q_{(t,2)} = \exp\left(-\int_0^t \beta_{(x,2)}\,\mathrm{d}x\right) = \exp\left[-\lambda q\int_0^t z + (1-z)\left(1 - e^{-\lambda q(\tau-x)}\right)^2\mathrm{d}x\right]$$

$$= \exp\left\{-\lambda q t + (1-z)\left[2\left(e^{-\lambda q(\tau-t)} - e^{-\lambda q\tau}\right) - \frac{1}{2}\left(e^{-2\lambda q(\tau-t)} - e^{-\lambda q\tau}\right)\right]\right\}$$

So $\dot{A}_t = (\lambda q)^2(1-z)\left[z - (1-z)\left(1 - e^{-\lambda q(\tau-t)}\right)\left(3 + e^{-\lambda q(\tau-t)}\right)\right]$

$$\times \exp\left\{\lambda q(t - 2\tau) + (1-z)\left[2\left(e^{-\lambda q(\tau-t)} - e^{-\lambda q\tau}\right) - \frac{1}{2}\left(e^{-2\lambda q(\tau-t)} - e^{-\lambda q\tau}\right)\right]\right\} \quad (2.B.7)$$

The first term in square brackets determines the sign of the profile slopes. It is positive if condition (2.31) holds, i.e., for

$$z \ge \zeta_{(t,2)} = 1 - \left(1 + S_{(t,2)}^{\text{bid}}\right)^{-2} = 1 - \left(4 - 4e^{-\lambda q(\tau-t)} + e^{-2\lambda q(\tau-t)}\right)^{-1} \ \forall t \in [0, \tau] \quad (2.B.8)$$

A strict inequality implies a strictly positive slope for any $\zeta_{(t,2)} < z < 1$. For $t = 0$ and $\lambda q\tau = 0.126$, the z lower bound is $\zeta_{(0,2)} = 0.2$.

2.C   Supplementary theory

This appendix adds formal details behind the infinitesimal generator and related results used in the paper to characterize dynamics. 2.C.1 provides a full definition. 2.C.2 uses it for a formal proof of Eq. (ODE-S) and offers an alternative derivation of Eq. (REC-S). 2.C.3 presents full details on the differential equations that determine the transition probabilities. 2.C.4 and 2.C.5 derive jump variance and martingale equivalences.

2.C.1   Infinitesimal generator

The Poisson count process B_t equals the number of bids by date t. Recall our state variable is $g_t \triangleq g_0 - B_t$. Denoting the change in g_t over an infinitesimal interval of length dt by $dg_t \triangleq g_{t+dt} - g_t$, the stochastic differential equation (SDE) for g_t is $dg_t = (-1) \times dB_t$ (SDE-g). Given that the instantaneous probability of a bid is $\beta_{(t,g)}\, dt$ and that of two or more bids occurring simultaneously is negligible, $\mathbb{P}((-dg_t) = 0|(t,g)) = 1 - \beta_{(t,g)}\, dt + O(dt)^2, \mathbb{P}((-dg_t) = 1|(t,g)) = \beta_{(t,g)}\, dt + O(dt)^2$ and other values can be ignored.[14] So for any process $Y_{(t,g_t)}$ adapted to g_t, Itô's formula applied to the Poisson jump process (SDE-g) yields the jump-diffusion

$$dY_{(t,g_t)} = \dot{Y}_{(t,g_t)}\, dt + \left(Y_{(t,g_t-1)} - Y_{(t,g_t)}\right)(-dg_t) \qquad \text{(SDE-}Y\text{)}$$

The *infinitesimal-generator* of $Y_{(t,g_t)}$ is its limit expected rate of change and satisfies

$$\mathcal{L}^Y_{(t,g)} \triangleq \lim_{dt \downarrow 0} \frac{\mathbb{E}_{(t,g)}\left(Y_{(t+dt,g_{t+dt})}\right) - Y_{(t,g)}}{dt} = \dot{Y}_{(t,g)} + \beta_{(t,g)}\Delta Y_{(t,g)} \qquad \text{(GEN)}$$

Recall that $\Delta Y_{(t,g)} \triangleq Y_{(t,g-1)} - Y_{(t,g)}$, distinct from $dY_{(t,g_t)}$.

Unlike total derivatives in standard calculus, the generator $\mathcal{L}$ satisfies the following adapted chain rule: for any twice differentiable function $h(\cdot)$ of $Y_{(t,g_t)}$ with $h_Y \triangleq \partial h(Y)/\partial Y$,

$$\mathcal{L}^{h(Y)}_{(t,g)} = h_Y \mathcal{L}^Y_{(t,g)} + \beta_{(t,g)}\left(\Delta h(Y_{(t,g)}) - h_Y \Delta Y_{(t,g)}\right) \qquad \text{(GEN-}h(Y)\text{)}$$

To see this, substitute $h(Y_{(t,g)})$ for $Y_{(t,g)}$ in Eq. (GEN), noting that $\frac{\partial}{\partial t}h\left(Y_{(t,g)}\right) = h_Y \dot{Y}_{(t,g)}$ and substitute for $\dot{Y}_{(t,g)}$ from Eq. (GEN).

---

[14]We use big-O notation, writing $y(t) = O(h(t))$ as $t \to 0$ if $\exists M, \epsilon : |y(t)| \le Mh(t)$ for all $|t| \le \epsilon$.

2.C.2   Alternative derivation of the success probability recursion

As $S_{(t,g_t)}$ is a martingale, $\mathcal{L}^S_{(t,g)} \equiv 0$ (see 2.C.5), so $\dot{S}_{(t,g)} = -\beta_{(t,g)}\left(S_{(t,g-1)} - S_{(t,g)}\right)$. Since $\beta_{(t,g)}$ is determined by $S_{(t,g-1)}$, we solve this first-order, linear, non-homogeneous differential equation (Eq. (ODE-S)) for $S_{(t,g)}$ given $S_{(t,g-1)}$: the integrating factor is $n_T^{(t,g)} = \exp\left(-\int_t^T \beta_{(x,g)}\,\mathrm{d}x\right)$ (Eq. (2.7)'s probability of no bid on (t,T) given gap $g_t = g$),

$$D_T\left(n_T^{(t,g)} S_{(T,g)}\right) = -n_T^{(t,g)}\beta_{(T,g)}S_{(T,g)} + n_T^{(t,g)}\dot{S}_{(T,g)} = -n_T^{(t,g)}\beta_{(T,g)}S_{(T,g-1)}$$

Integrating, using $n_t^{(t,g)} = 1$ and $S_{(\tau,g)} = 0$ for $g \geq 1$, (ODE-S)'s unique solution is

$$S_{(t,g)} = \int_t^\tau n_T^{(t,g)}\beta_{(T,g)}S_{(T,g-1)}\,\mathrm{d}T$$

$S_{(t,g)} \equiv 1 \,\forall g \leq 0$ completes this recursive solution as alternative derivation of (REC-S).

2.C.3   State transition probabilities

As with the alternative derivation of Eq. (ODE-S) provided in Section 2.3.2, we begin by proving Eq. (REC-Q) via the adjoint to Eq. (ODE-Q). This adjoint is called the Kolmogorov backward equation because it fixes the target state (t,g) and solves for Q by integrating backwards to the earlier state (t',g'). To derive the adjoint, we sum the probabilities of reaching (t,g) from (t',g') via $(t'+\mathrm{d}t',g')$ with no bid on infinitesimal interval $(t',t'+\mathrm{d}t')$ and the alternative path via $(t'+\mathrm{d}t',g'-1)$ with one bid on $(t',t'+\mathrm{d}t')$:

$$Q_{(t,g)}^{(t',g')} = \left(1 - \beta_{(t',g')}\,\mathrm{d}t'\right)Q_{(t,g)}^{(t'+\mathrm{d}t',g')} + \left(\beta_{(t',g')}\,\mathrm{d}t'\right)Q_{(t,g)}^{(t'+\mathrm{d}t',g'-1)} + O\left(\mathrm{d}t'\right)^2$$

$$\text{So}\ \lim_{\mathrm{d}t'\to 0}\frac{Q_{(t,g)}^{(t'+\mathrm{d}t',g')} - Q_{(t,g)}^{(t',g')}}{\mathrm{d}t'} = -\beta_{(t',g')}\left(Q_{(t,g)}^{(t',g'-1)} - Q_{(t,g)}^{(t',g')}\right)$$

$$\text{So}\ D_{t'}\left(Q_{(t,g)}^{(t',g')}\right) = -\beta_{(t',g')}\left(Q_{(t,g)}^{(t',g'-1)} - Q_{(t,g)}^{(t',g')}\right) \qquad \text{(ODE-}Q\text{-backward)}$$

The integrating factor to solve (ODE-Q-backward) is $n_T^{(t',g')}$. For $g \leq g'-1$, this generates the recursive solution,[15]

$$Q_{(t,g)}^{(t',g')} = \beta_{(T,g')}\int_{t'}^t n_T^{(t',g')}Q_{(t,g)}^{(T,g'-1)}\,\mathrm{d}T \equiv \mathbb{E}_{(t',g')}\left[Q_{(t,g)}^{(T_{g_0-g'+1},g'-1)}\right]$$

which is precisely Eq. (REC-Q). For $g = g'$, given that $Q_{(t,g)}^{(t,g)} = 1$ and $Q_{(t,g)}^{(t,g'')} = 0$ for all $g'' < g$, the solution is Eq. (2.24), confirming the overall solution derived in the text.

---

[15]Since T_n is the n'th stopping time, the next bid after (t',g') occurs at $T_{g_0-g'+1}$.

To derive the Kolmogorov forward equation that we use in Lemma 2.3, we instead fix the initial state $(t', g') = (0, g_0)$, restricting away from the general case, and look forward to assess the probability of marginal shifts in the later state (t, g). A campaign reaches $(t + dt, g)$ via either a bid from just prior state $(t, g + 1)$ or via (t, g) with no intervening bid on the interval $[t, t + dt)$. So

$$Q_{(t+dt,g)} = Q_{(t,g)}(1 - \beta_{(t,g)}\,dt) + Q_{((t,g+1))}\beta_{(t,g+1)}\,dt$$

Using $Q_{(t+dt,g)} - Q_{(t,g)} = \dot{Q}_{(t,g)}\,dt$ and dividing by dt and taking limits as $dt \to 0$, this gives the differential equation Eq. (ODE-Q). We suppress the "-forward" qualification since the main paper only needs this single variant.

2.C.4   Jump variance

Using Eq. (SDE-Y) and neglecting terms of order $O(dt^2)$,

$$\nu_{(t,g)}^Y \triangleq \lim_{dt\downarrow 0} \frac{\mathbb{V}_{(t,g)}\left(Y_{(t+dt,g_{t+dt})} - Y_{(t,g)}\right)}{dt} \qquad = \lim_{dt\downarrow 0} \frac{\mathbb{V}_{(t,g)}\left(dY_{(t,g)}\right)}{dt}$$

$$= \lim_{dt\downarrow 0} \frac{\mathbb{E}_{(t,g)}\left[\left(Y_{(t+dt,g-1)} - Y_{(t,g)}\right)\left(-dg_t\right)\right]^2}{dt} \qquad = \beta_{(t,g)}\left(\Delta Y_{(t,g)}\right)^2$$

since the instantaneous variance of the underlying Poisson process g_t equals its intensity $\beta_{(t,g)}$; each bid causes $Y_{(t,g)}$ to jump by $\Delta Y_{(t,g)}$; squaring that scale factor gives the result.

2.C.5   Martingale equivalences

Using the transition probabilities determined in Lemma 2.2, the expected generator at time x given an initial state (t', g') can be expressed as

$$\mathbb{E}_{(t',g')}\left(\mathcal{L}_{(x,g_x)}^Y\right) = \sum_{g \leq g'} Q_{(x,g)}^{(t',g')}\mathcal{L}_{(x,g)}^Y$$

To see that the generator of a martingale process Y_{t,g_t} must be identically zero and always positive for a submartingale, always negative for a supermartingale, notice that

$$\frac{d}{dt}\left(\int_{t'}^t \sum_{g \leq g'} Q_{(x,g)}^{(t',g')}\mathcal{L}_{(x,g)}^Y\,dx\right)\Bigg|_{t'=t} = \sum_{g \leq g'} Q_{(t',g)}^{(t',g')}\mathcal{L}_{(t',g)}^Y = \mathcal{L}_{(t',g')}^Y$$

So the martingale property is violated if there exists a (t, g) such that $\mathcal{L}_{(t,g)}^Y \neq 0$. The proof for supermartingales and submartingales is also immediate from this identity. To prove the converse implication and complete the proof of the equivalences, it suffices to

apply Dynkin's Theorem (see e.g., Theorem 1.24 in Øksendal and Sulem (2007)) which states that for any $t' \leq t$: $\mathbb{E}_{(t',g')}(Y_{(t,g_t)}) - Y_{(t',g')} = \int_{t'}^{t} \mathbb{E}_{(t',g')}\left(\mathcal{L}_{x,g_x}^{Y}\right) dx$.

2.D Supplementary Figures
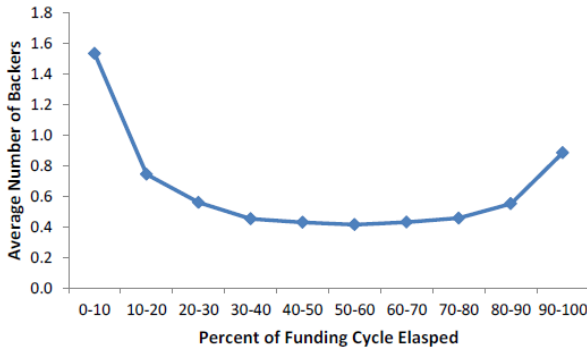
(a) Kuppuswamy and Bayus (2015)
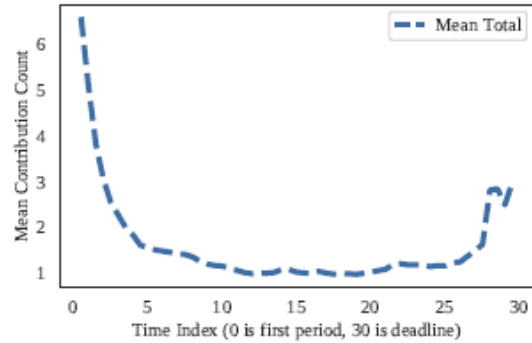
(b) Deb et al. (2021)

(c) Crosetto and Regner (2018)

Figure 2.D.1: All panels plot number of bids per time interval. Data sources are Kickstarter (c,b) and Startnext (c). (b) restricts to 30 day duration campaigns while (c,c) measure time as a percentage of duration. (c) distinguishes campaigns based on funding progress relative to the average path of the funding ratio (cumulative bids divided by goal). "ontrack" projects deviate by 5% at most form the average.
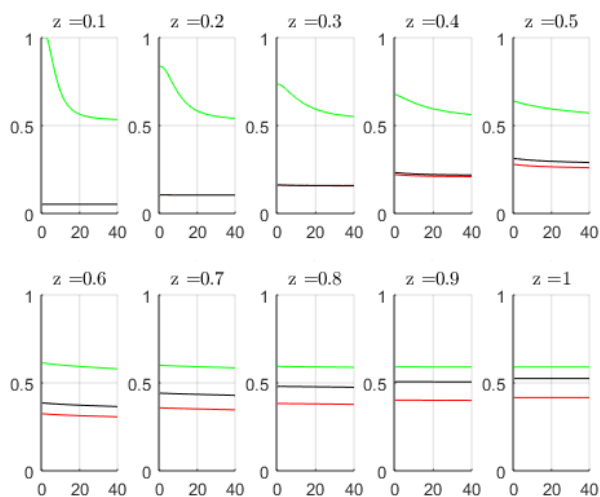
(a) Linear CDF ($\rho = 1$)
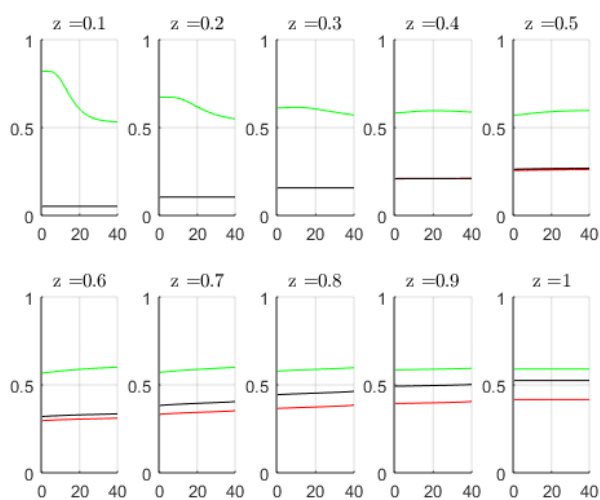
(b) Quadratic CDF ($\rho = 2$)
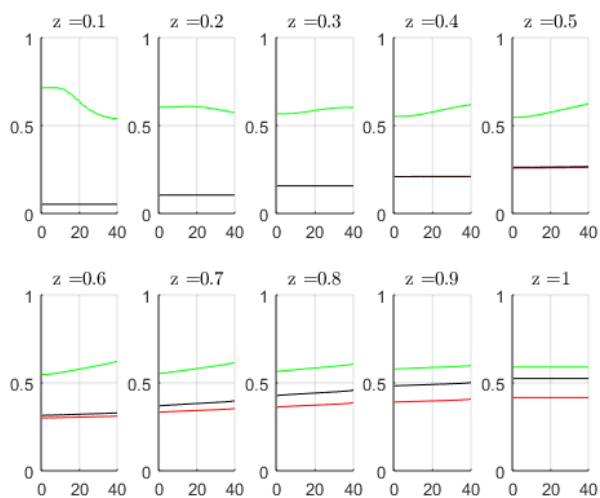
(c) Cubic CDF ($\rho = 3$)

Figure 2.D.2: Effect of z on bid profiles $\textcolor{green}{\longrightarrow} A_t^{\mathcal{S}}$ $\longrightarrow A_t$ $\textcolor{red}{\longrightarrow} A_t^{\mathcal{F}}$ with a power-ρ distribution function $F(c) = z + (1 - z)(c/q)^\rho$ given $g_0 = 20$ and $(\tau, \lambda, q) = (40, 0.7, 0.75)$ ◆

Page intentionally left blank.

Chapter 3

Crowdfunding with Endogenously Timed Moves

(Joint with Matthew Ellman)

3.1   Introduction

Chapter 2 studies bidding dynamics and crowdfunding outcomes assuming an exogenous temporal order of bidder moves. In reality, bidders rarely control exactly when they hear about a specific campaign, but once they become aware, they may be able to choose to decide later. By delaying a bidder can learn about the bidding choices of others who choose to move in the interim. Continuous time delay, in the sense of following a campaign over time, requires prolonged attention that is too costly for most bidders. However, many platforms, including Kickstarter, offer a "Remind me" button that makes it easy for bidders to revisit once, shortly before the deadline. An effective revisit still involves duplication of some thinking costs – the cost of registering basic project characteristics into the bidder's short-term memory ready for deciding what to do – but in principle that cost can be very small. So this paper extends the exogenous sequentiality model of Chapter 2, now denoted Exo, by allowing this type of endogeneity in the timing of moves: bidders can delay their decisions till just before the campaign deadline at no cost. We investigate what these delay options imply for bidder surplus, success rates and total welfare, compared to Exo and a simultaneous move benchmark denoted Sim. We also analyze the implications for bidding dynamics.

The following core features of the Exo model continue to apply. Costs of inspecting the entrepreneur's campaign continue to play a central role in the analysis. Each bidder learns the campaign's current state for free on arrival. The observable state is summarized by the date t and funding gap g_t, as well as the product price p and deadline τ. Campaigns publicize the amount of funds cumulated so far and the funding threshold from which bidders can immediately learn the funding gap by subtraction. The only substantative

learning cost for a bidder is the cost of inspecting to discover his individual valuation which is a private value. The new question is whether and when bidders want to delay this costly inspection and bidding decision. In EXO, they had no such option. In ENDO, they can. The number who delay and revisit the campaign at its end changes the strategic problem of every bidder. In EXO, average pivotality is decreasing in time because later bidders have fewer successor bidders to influence but in ENDO, each bidder can influence his predecessors if they chose to delay. A key difficulty is that bidders cannot observe how many predecessors arrived nor how many of them chose to delay. Instead bidders use the campaign's state variables to estimate how many bidders will revisit the campaign at its deadline.

Bidders who delay can respond to news about bidding by those who act on arrival but delay options open the door to free-riding in the form of excessive delay. This can lower success rates and welfare relative to a simultaneous benchmark (SIM) and exogenous sequentiality (EXO). On the other hand, endogeneity (ENDO) potentially raises welfare because low cost types endogenously move earlier. Endogenous timing can even raise campaign success rates when there are enough high cost types who do not inspect in SIM, who rarely inspect in EXO but who are always willing to delay in ENDO.

We fully characterize all equilibria in the special case of a threshold of two. Bidders tradeoff a higher chance of getting the good by inspecting on arrival agains cost-saving via free-riding by delaying. With a two bid threshold, their rankings of these two options, to inspect or delay, turn out to be time-independent.

If inspection costs are homogenous or binary, there is a continuum of pure and mixed strategy equilibria. These equilibria all share the same expected number of potential delayers; that is, bidders who delay if the gap is still equal to the original threshold of two. Only the equilibrium in stationary strategies is robust to inspection cost perturbations. When the distribution of inspection costs is continuous, the model delivers a unique prediction in stationary strategies.

When inspection costs are homogeneous, it is easy to compare the three crowdfunding models. Given a threshold two campaign, bidders' delay strategies generate the same bidder surplus in each of the three models. The success rate, by contrast, is higher in SIM than in the other two models. The reason is that simultaneity prevents bidder free-riding and prevents bidders getting discouraged from learning bad news about their predecessors. As a result, SIM is the best design option for the entrepreneur when bidders are homogeneous. This advantage of SIM is greater at higher thresholds because higher thresholds make free-riding more harmful.

Welfare comparisons between SIM, EXO and ENDO are less clear-cut when inspection costs are heterogenous. If a bidder type abstains (avoids inspecting and bidding) in SIM and usually abstains in EXO, then ENDO provides the highest surplus for that bidder type. This is because such bidders are at least willing to delay and may strictly benefit

from the option to inspect at the last moment. In addition, moderate cost types who would inspect on arrival at certain dates in Exo can gain information rents by delaying in Endo and this may raise overall bidder surplus.

In the two-bid threshold setting, the expected surplus of low cost types who continue to inspect is unaffected and bidder surplus rises if the overall level of free-riding is not too high. As always, Sim maximizes the success rate if all bidders are willing to inspect in Sim but when the probability measure of high cost types is high, having all bidders always inspect is not sustainable even under simultaneous moves. Action sequencing and endogeneity then has the advantage that high-cost types can move after low-cost types. This is an advantage because they then observe when success becomes more likely and at least sometimes inspect. So long as there is not too much free-riding among low cost types, Endo then raises success rates relative to Exo. Exo can also improve on Sim because high cost types sometimes exogenously arrive after low cost types, albeit less often than in Endo.

In general, endogenous timing has greater benefits in adverse crowdfunding environments, defined by low bidder priors about product value and high inspection costs. In those settings, endogeneity is important because delay by high cost types raises the chance of taking full advantage of the relatively few zero and very low cost types who do arrive. Conversely, simultaneity helps the most when project fundamentals are strong and inspection costs are concentrated on low but positive values. In that case, abstention is not a concern but free-riding from bidder delay may still occur and reduce success rates.

Returning to welfare considerations, we also solve for the constrained first-best in which a planner decides whether bidders inspect or delay or abstain. We find that the optimal amount of delay is positive but below that in Endo, so Sim has too little delay but Endo has too much. The distortion in Endo is caused by the negative externality imposed by free-riders on each other.

Finally, we study the bid profiles generated by the equilibria of Endo as compared to those in Exo. In practice, most crowdfunding campaigns are advertised in advance of official initiation, we extend Endo by adding an initial population of pre-arriving bidders who act as soon as the campaign starts. We denote the resulting extended model by Endo-Pre. The obvious addition to Exo is that the These two features clearly permit the bid profile to have initial and final spikes. We prove that those spikes do indeed materialize in equilibrium and thereby generate a more pronounced U-shape or bimodal profile than in Exo.

The focal stationary equilibria of Endo induce increasing dynamics in the intermediate stages of a campaign. This results in a bimodal profile when pre-arrivals are present. A final spike in the bidding profile always arises except as one possible equilibrium in the knife-edge case where costs are perfectly homogenous; that is the equilibrium that is profile-equivalent to Exo; i.e., such that all bidders delay in Endo at the same dates at

which they would abstain from inspecting and bidding in Exo.

Related literature

The most closely related papers are Liu (2020) and Deb et al. (2021). Relative to these contemporaneous papers, our paper is unique in that we have multiple actors moving endogenously in a setting with conflicting interests among bidders.

Liu (2020) has multiple endogenously moving players in a common value crowdfunding, but she assumes that bidders have no cost in placing their bids and identical ex-post payoffs. Her appendix shows that bidders sort with optimists first in a perfect descending signal timing that fully reveals their signals. This relies on the fact that bidder incentives are perfectly aligned. In our setting instead, bidders have a common interest in the success of the campaign but a bidder who likes the good wants others to inspect and bid even if their inspection costs are too high and they do not like the good. Differences in costs and tastes create an incentive to free-ride. In terms of bid profiles, her model differs from our in that all bidders arrive at the start. So she predicts a time-schedule of bidder action that is arbitrary except for ranking bidders according to their signal precision.

In Deb et al. (2021), the timing of action is endogenous in that a donor continuously follows the crowdfunding campaign and decides how much to pledge at any point in time. This donor can pledge multiple times. Besides the donor, also bidders, called "buyers", join the campaign at a homogeneous Poisson rate and decide whether to bid or not, on arrival, as in Exo. They have no inspection cost for buyers but assume a cost of placing a bid. This has similar effects. As they acknowledge, the buyers could donate like the donor but they assume a single donor actor because the analysis is otherwise intractable. Our contribution is to allow for multiple actors who can delay in a specific way. The delay is specific but plausible as motivated by the prevalent Remind-me button.

These authors also study optimal design from the point of view of bidders, the platform, and the strategic donor by characterizing welfare properties of equilibria and presenting alternative campaign design possibilities to their full information baseline. Their "No information" environment coincides with our Sim model when the donor has no wealth. We also ask whether design options are platform or bidder optimal by analysing whether they imply maximal success rates and bidder surplus. We study optimal design in terms of total welfare and provide a fuller characterization. With homogenous costs, we find that Sim is both bidder and platform optimal; hence it is also optimal in terms of aggregate welfare. Deb and coauthors find an analogous result when the donor's success motivation is low. In contrast, with heterogenous costs, we find that Endo is always bidder-optimal and sometimes also platform-optimal, in the sense of maximizing the success rate. For instance, Endo maximizes success when the cost distribution has a thick upper tail.

To rule-out unstable equilibrium predictions, we define the equilibrium refinement "strict-dominance-elimination-then-weak-dominance-elimination", abbreviated as 1S1W. This eliminates weakly dominated strategies after first eliminating the strictly dominated ones. Our refinement is intuitively implied by Stelten's perfect equilibrium and Myerson's proper equilibrium (see Selten (1975), Myerson (1978), and Van Damme (1991)), but differs from these standard refinements that do not apply to our context due to the Poisson arrival of bidders that implies a potentially infinite number of players and periods; SIM is a standard Poisson game as defined in Myerson (1998). An advantage of 1S1W is that it is easy to verify.

Our study is more distantly related to a range of models in continuous time with endogenous move order among which two on bargaining and irreversible investment are the closest. Zhang (1997) studies investment cascades over a finite horizon in continuous time. In his equilibrium, as in Liu (2020), bidders with positive signals plan to act in descending order of signal precision if no bid has yet been pledged, but given that waiting is costly, once the first bid is pledged by the bidder with the most precise signal, all other bidders pledge immediately since waiting is costly and they cannot acquire more precise information. The timing of actions after the first pledge would be arbitrary in the absence of the waiting cost.

3.2 Model

This section introduces an action delay option into the crowdfunding model presented in Chapter 2, which we call EXO since the timing of action is purely exogenous. We offer a recap EXO and then present the ENDO model, so named because action delays make sequencing endogenous. We also introduce SIM, a variant of the crowdfunding game where all bidders act at the same time. SIM provides the natural static benchmark.

All model variants describe an All-or-Nothing (AoN) crowdfunding campaign. The entrepreneur makes an open call to raise enough bids at a price p to cover a bidder threshold of g_0 bids during a time period of duration τ. If she succeeds (event $\mathcal{S}$), bidders pay their bids and the entrepreneur launches production paying a sunk investment cost and delivering the crowdfunding product to those who bid. In the opposite case where the campaign fails (event $\mathcal{F}$), production does not occur and bids are not paid. For each $t \in [0, \tau_+]$, the funding gap g_t counts the additional bids required to reach the funding threshold. It starts at g_0 and makes a unit drop when a bid occurs. At time τ_+, bids are added up to decide the outcome of the campaign: success $\mathcal{S} \triangleq \{g_{\tau_+} \leq 0\}$ or fail $\mathcal{F} \triangleq \{g_{\tau_+} > 0\}$.

3.2.1   Recap of exogenous sequentiality (Exo)

Bidders arrive or discover the campaign according to a homogenous Poisson process with rate λ over $t \in [0, \tau)$ and choose what to do on arrival.[1] We refer to a specific bidder by his arrival date t.

Since bidders are risk-neutral, their payoffs are linear in probabilities. A bidder valuation for the crowdfunding product can be either $v_t = 0$ or $v_t = v > 0$ with prior probability q on v. Bidders do not know v_t upon arrival but can pay an inspection cost c_t to learn it and bid conditional on learning $v_t = v$. After normalizing the price to $p = v - 1$ so that bidders who value the good enjoy one unit of utility from purchasing the good *net* of this price, a bidder obtains a payoff of $u^{\mathbf{A}} = 0$ by **A**void bidding and avoid inspecting (**A**), while making a **B**lind bid without prior inspection (**B**), he obtains

$$u^{\mathbf{B}} = \mathbb{1}_{g_{\tau_+} \leq 0} \left(v_t - (v - 1) \right) \tag{3.1}$$

In words, if the campaign is successful, the bidder gets the value for the crowdfunding product and pays its price.

The payoff (3.1) is clearly negative if $v_t = 0$; that is, bidders who know they have the zero valuation never want to buy. So bidders face a risk of a mistaken purchase. They can avoid that risk by first paying the inspection cost c_t to **C**heck that $v_t = v$ and bid only in that case (**C**). The payoff associated with this strategy is

$$u^{\mathbf{C}} = \mathbb{1}_{g_{\tau_+} \leq 0} \left(v_t - (v - 1) \right)^+ - c_t \tag{3.2}$$

so a bidder playing **C** only bids for the good when he likes it and then always gains from a crowdfunding success but his inspection cost is sunk independent of the crowdfunding outcome. We assume that inspection costs are i.i.d. draws from the cumulative distribution function (CDF) $F(\cdot)$ with support taken to be a subset of the $[0, q]$ interval so that inspection gives a positive expected payoff to all cost types if success is certain. As we show in Chapter 2, this restriction is without loss of generality since all bidders with $c_t \leq 0$ behave as if they had $c_t = 0$ while those with $c_t > q$ never inspect. Hence assuming a positive density over that region is equivalent to reducing the arrival rate of bidders that do inspect.[2]

We make the No-Blind-Bidding assumption that q and v are low enough to discourage bidders from playing **B**. Specifically, we assume $q < 1 - 1/v$ (NBB) so that strategy **B** strictly dominated by **A**. This way a bid always follows an inspection decision, so that

---

[1]Writing the interval as $[0, \tau]$ is essentially identical since the probability of arrival at any instant, including that final instant, is zero. Ruling out arrivals at τ simplifies the exposition since the bidder arriving exactly at the return date $t = \tau$ would not have the option to delay that all other bidders have and his bid could not influence the returning delayers.

[2]If F is discrete, our formulation is equivalent to assuming that each cost type arrives according to an independent Poisson process.

"inspect" always refers to inspect and bid if v_t is found to equal v.

3.2.2 Endogenous sequentiality (Endo)

The key addition to Chapter 2 is the introduction of the strategy "**D**elay" (**D**) that allows bidders to postpone their action at time τ, when the campaign is about to end. Specifying a full strategy requires bidders to decide what sub-strategy to take when they come back at τ for any given gap g_τ that they will observe. Given that the only possible undominated actions when returning are either to abstain or to check the campaign, the strategy $\mathbf{D}\overline{g}$ such that bidders delay and inspect only if $g_\tau \leq \overline{g}$ specifies a full contingent plan. A bidder's payoff from choosing this strategy is

$$u^{\mathbf{D}\overline{g}} = \mathbb{1}_{g_\tau \leq \overline{g}} \left(\mathbb{1}_{g_{\tau_+} \leq 0} \left(v_t - (v-1) \right)^+ - c_t \right) \tag{3.3}$$

The difficulty of dealing with the large strategy space obtained by combining all the contingent responses to g_τ becomes an obstacle for large g_0 campaigns. It is therefore valuable to restrict attention to the single delay strategy $\mathbf{D}\overline{1}$. We denote this $\mathbf{D} \equiv \mathbf{D}\overline{1}$ from now on. A bidder who chooses $\mathbf{D}$ inspects only when success is certain and bids only after learning he likes the good. We verify that relaxing this restriction does not change the equilibrium in our solutions: When $g_0 = 2$, waiting and bidding with $g_\tau = 2$ is ruled-out by 1S1W.

3.2.3 Simultaneous moves benchmark (Sim)

To assess whether removing all forms of sequentiality might be optimal, we also solve a game where all bidders move simultaneously and the number of bidder arrivals is Poisson distributed with parameter $\lambda\tau$. This "equivalent simultaneous game" SIM is a standard Poisson game (see Myerson (1998)).

The simultaneous setting is strategically equivalent to a crowdfunding campaign with no updating of information over time, such as a campaign where all bidders join at $t = \tau$ or with g_t not observable throughout the whole campaign duration. In these circumstances, the ENDO and EXO distinction of model variants (later, bidder types) is irrelevant. Since there is no substrategy $\mathbf{D}$ in SIM, the only possible substrategies after applying assumption (NBB) are $\mathbf{A}$ and $\mathbf{C}$ as in EXO.

3.2.4 Equilibrium concept

The equilibrium concept we employ is Perfect Bayesian Equilibrium refined by two additional criteria: Pareto Optimality and the novel refinement "strict-dominance-elimination-then-weak-dominance-elimination", denoted *1S1W* for brevity, such that only strategies

that are weakly undominated among those that survive one round of elimination of strictly dominated strategies can constitute an equilibrium. The equilibrium concept can be abbreviated as 1S1W-P-PBE.

We also impose Pareto Optimality to rule the possible coordination failure where each bidder plays **A** believing that all others do the same. This can arise with simultaneous play, as when all delayers return in ENDO and in the single bidding round in SIM. Our *1S1W* makes equilibria robust to small errors (trembles) in bidders' implementation of their planned strategies or in their assessment of others' equilibrium play. Applied to ENDO, this rules out substrategy **A**; rather than abstaining, bidders wait until the end in the hope that other bidders pledge. **A** is eliminated by 1S1W precisely because, after **C** eliminates by strictly dominating **A** for $g_t \leq 1$, **D** eliminates **A** by weak domination for $g_t > 1$.

One might think that tie-breaking rules such as bidders having a lexicographic preference for pro-social behavior (or equivalently, for promoting production or social welfare) could refine nicely the equilibrium set. With this lexicographic preference **C** wins in a tie with **A** or **D**, but we avoid assuming so since discontinuous utility functions create problems for equilibrium refinement in Poisson settings. We make an exception to this rule tho for strategy choices that are irrelevant, in the sense that either the probability of a bidder arrival is zero, or the bidder strategy has no impact on other players' payoffs. In these circumstances we do tie-break in favour of **C**.

As a result of our refinements and the (NBB) assumption, given common knowledge of prior distributions through $(\lambda, q, F(.))$ and observable campaign features $(g_0, \tau, g_t$ in EXO and ENDO; only g_0 and τ in SIM), equilibrium strategies are mappings $a : (t, g_t, c_t) \rightarrow \{\mathbf{C}, \mathbf{D}\}$ in ENDO; $a : (t, g_t, c_t) \rightarrow \{\mathbf{A}, \mathbf{C}\}$ in EXO; $a : c_t \rightarrow \{\mathbf{A}, \mathbf{C}\}$ in SIM. We define the pair (t, g) as the campaign's *funding state* so $g_t = g$ and we denote the strategy played by a generic bidder at a particular funding state by $a_{(t,g)}$.

In EXO, our assumption that bidders observe the gap at the date they arrive implied that bidders knew all the history of play that can affect them; they do not care whether a no bid episode results from a non-arrival, an arrival who chose **A** or an arrival who chose **C** but did not value the reward. As a result, the game is Markovian and has a unique equilibrium prediction.

In ENDO, we retain the assumption that the g_t is observed, but it no longer implies that the game has effectively perfect information about the past. The Poisson arrival process for bidders results in three potentially relevant state variables: the number of bidders who arrived and chose to bid, the number who arrived and decided definitively against bidding and the number who arrived and chose to delay. Of these, the first and last are relevant but only the first is observable. Even if all bidders observed the full history of bids up until the date at which they arrive or return, they would not discern the exact number of times actions **C** and **D** (and **A**) have been played since they all of

the above strategies can feature no bid. Notice that strategy **D** implies no bid but *if known* would be more encouraging for other bidders in terms of optimism about success. Communication among bidders would allow them to disclose their actions but we preclude such possibility.[3]

Arriving bidders at funding states (t, g) care about knowing the number of delayers that will return and inspect at the last moment of the campaign, but, due to this action unobservability, they only know that this number is Poisson distributed parameter $\lambda q \mu(\mathbf{D})$, where $\mu(\mathbf{D}) \triangleq \mathbb{E}\mu(\mathcal{T}^{\mathbf{D}})$ is the *expected* measure of dates $\mathcal{T}^{\mathbf{D}}$ at which, *if* a bidder arrives, that bidder plays **D**, given an equilibrium prescription. To define $\mathcal{T}^{\mathbf{D}}$ precisely, we denote the stopping time of the b'th bid by $T_{g_0 - b}, b \in \mathbb{N}$ and the set of dates at which **D** is played when the gap is *fixed* at a value of g as $\mathcal{T}^{\mathbf{D}|g} \triangleq \left\{ x \in [0, \tau) : a_{(x,g)} = \mathbf{D} \right\}$. When all stopping times of bids and inspection costs are realized, the set $\mathcal{T}^{\mathbf{D}}$ results from the union of the subsets of each of the $\mathcal{T}^{\mathbf{D}|g'}$ sets, for $g_0 \geq g' \geq 1$ that are crossed by the path of the gap. In formula, letting for convenience $T_0 \equiv 0$,

$$\mathcal{T}^{\mathbf{D}} = \bigcup_{b=1}^{g_0 - 1} \left\{ [T_{b-1}, T_b] \cap \mathcal{T}^{\mathbf{D}|g_0 - b + 1} \right\} \tag{3.4}$$

Bidders will have to estimate $\mu(\mathbf{D})$ based on the observed funding state and equilibrium prescriptions. Notice that the stopping times T_b are uncertain and bidder t only knows that $T_{g0-g} \leq t \leq T_{g0-g+1}$, with the second weak inequality holding as equality if bidder t chooses to bid. Therefore, since beliefs are not fully captured by observed funding states, the game lost its Markovian property and can now admit multiple equilibria. Note that even though arrivals are independent

Some further remarks are in order before proceeding: First, c_t does not matter for bidder t's beliefs about *other* bidders since inspection costs and arrival times are independent draws. However, the distribution of stopping times T_b are correlated, making the computation of $\mu(\mathbf{D})$ based on Eq. (3.4) too complicated in the general case for explicit solution. However, it is feasible to solve explicitly in simple cases as we show in the following sections. Second, if we were to allow generic $\mathbf{D}\overline{g}$ strategies, we would also have to take into account beliefs about the composition of the crowd of delayers. When restricting to **D**, those beliefs do not matter because success is certain when a delayer bids upon returning.

---

[3] Even the platform may not know if a bidder who arrived on a project webpage and did not bid was: (a) thinking seriously about bidding before learning that he did not like the reward enough, as represented by **C** plus the negative information signal represented by the low valuation $v_t = 0$, (b) thinking whether to bother inspecting and then deciding on **A** or **D**, or (c) simply visiting the project webpage and gazing absent-mindedly or thinking about something entirely different, which are equivalent to not arriving. Of course, the platform may have some additional information and bidders might signal their presence by sending messages but signals can be faked so, in the interests of realism, we allow for imperfect observability even though it complicates the analysis.

3.2.5   Welfare

An addition relative to Chapter 2 is that we provide welfare analysis. We define (expected) welfare W as the sum of two components. The first, bidder surplus, is given by $\lambda q V$ and results from multiplying the expected number of bidders $\lambda \tau$ by the expected utility of an arriving bidder $V \triangleq \mathbb{E}\left(u_{t_n}^a\right)$. This last expectation is conditional on the bidder arriving by τ and is computed with respect to the bidder arrival time t_n. The index n summarizes external analysts' knowledge about the bidder though the bidder only observes t_n and g_{t_n}. By default, we just denote a bidder arrival time by t which is what the bidder knows. Information sets (t, g_t, c_t) determine bidder strategies $a \in \{\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}\}$ that in turn determine the payoff $u_{t_n}^a$ of the n'th arriving bidder according to the above formulae (3.1), (3.2), and (3.3). Except in SIM where decisions depend only on g_0, an action's expected payoff will depend on g_{t_n}, which will itself depend on earlier arrival times and actions.

The second is the entrepreneur's value for a success R, which the entrepreneur earns with success probability $S_0 \triangleq \mathbb{P}(g_{\tau_+} \leq 0)$. The latter accounts for the surplus generated by crowdfunding completion, extracted by the entrepreneur in the form of reputational rents from after-market sales or from higher prominence on the crowdfunding platform. We do not take into account the implications of how revenues are split between crowdfunding and after-market though that would be interesting as future work.

$$W = \lambda \tau V + R S_0 \qquad\qquad V = \mathbb{E}\left(u_{t_n}^a\right), \quad S_0 \triangleq \mathbb{P}(g_{\tau_+} \leq 0) \qquad (3.5)$$

Notice that welfare maximization is achieved by maximizing the probability of crowdfunding success if R is sufficiently high. For this reason, the success rate is an important welfare metric. To facilitate the comparison of welfare values reached under the different assumptions we make on the timing of bidders' moves, we denote welfare by W^{SIM}, W^{EXO}, W^{ENDO} for the respective specific models. We adopt the same naming convention for V and S_0.

3.3   Results for benchmark structures

In this section, we present two generic benchmarks: SIM and EXO, that we analyze under the assumption that inspection costs homogenous (F is a unit-step function with step at c). SIM is elementary, so we solve its equilibrium and the implied welfare for a generic g_0. We already solve EXO in Chapter 2, so here we just carry the solution from there. In both cases, we solve welfare (bidder surplus and success rate) explicitly for $g_0 = 2$ since it will allow us to compare it with our results for ENDO. We deal with heterogenous inspection costs in benchmarks in Section 3.5 on optimal design.

3.3.1 Simultaneous game

Since SIM is an important benchmark and its equilibria are easier to solve than in the other model variants, we start from there. The aspect that makes it elementary compared to EXO and ENDO is that bidders do not take into account any dynamically evolving information; they decide to play **C** or **A** based solely on g_0 and their inspection cost c_t. Specifically, the expected utility that a bidder obtains by playing **C** is given by

$$U^{\mathbf{C}} = \mathbb{E}_{(0,g_0)}\left(u^{\mathbf{C}}\right) = qS^{\text{bid}} - c_t \tag{3.6}$$

where S^{bid} is the bid-conditional estimate of the success probability.

Equilibria are readily characterized based on Eq. (3.6), which shows that they have a cutoff structure: since S^{bid} is independent of c_t by the assumption of i.i.d. inspection costs, the expected utility in Eq. (3.6) is linear in c_t. It follows that equilibria are threshold values $\hat{c}$ such that bidders play **C** for $c \leq \hat{c}$ and **A** otherwise, so that the probability of a bidder inspecting, $qF(\hat{c})$, combined with the fact that the bidder population is Poisson distributed with parameter $\lambda\tau$ gives a bid-conditional success rate of

$$S^{\text{bid}} = \sigma_{g_0-1}\left(\tau F(\hat{c})\right) \triangleq 1 - \sum_{b=0}^{g_0-2} \frac{(\lambda q\tau F(\hat{c}))^b}{b!} e^{-\lambda q\tau F(\hat{c})} \tag{3.7}$$

where $\sigma_{g_0-1}(x)$ indicates the Poisson probability of at least $g_0 - 1$ bids from bidders arriving over a set of dates with measure (or expected measure) $x \geq 0$.

There can potentially more than a single threshold consistent with equilibrium, yet equilibrium selection based on Pareto optimality achieves uniqueness, with the unique Pareto optimal equilibrium given by the highest among those. Notice that $\hat{c} = 0$ is always a valid PBE but, according to the Pareto criterion, we discard it whenever there is another valid (and undominated) PBE with $\hat{c}' > 0$.

In the baseline of homogenous inspection costs $c_t = c$ we a unique (PO) equilibrium in which either all bidders choose **C** as long as each of them obtains a positive payoff from doing so; or all play **A**. When equilibrium prescribes bidding we have $c \leq \hat{c}$, implying $F(\hat{c}) = 1$. So bidders do not deviate to **A** as long as $U^{\mathbf{C}} \geq 0$; that is, $\sigma_{g_0-1}(\tau) \geq c/q$. If this condition holds, bidder surplus and success rate are given by

$$U^{\mathbf{C}} = V = q\sigma_{g_0-1}(\tau) - c \qquad\qquad S_0 = \sigma_{g_0}(\tau)$$

From here, determining welfare is straightforward as it results from combining (and scaling) the previous V and S_0 values. Of course, for $c > q\sigma_{g-1}(\tau)$ equilibrium prescribes **A** so all welfare quantities are zero. This concludes the analysis for the homogenous case in SIM.

In the preceding analysis, we characterized equilibrium using a cutoff $\hat{c}$ on c although

we could as well define a cutoff $\hat{\tau}(\lambda, c)$ on the minimal duration τ required for bidder to be confident enough that a success will occur after they bid. This second approach is more useful when comparing SIM with the other timing variants. So, as most of the following analysis will focus on $g_0 = 2$, we solve it explicitly for this case. For ease of exposition, we let

$$\sigma_1(x) \equiv \sigma(x) = 1 - e^{-\lambda q x} \tag{3.8}$$

denote the probability that at least one bid is pledged, which satisfies $\sigma_x(x) = \lambda q(1 - \sigma(x)) > 0, \sigma_{xx}(x) = -(\lambda q)^2(1 - \sigma(x)) < 0$ and the exponential property of closure under product $(1 - \sigma(x))^a(1 - \sigma(y))^b = 1 - \sigma(ax + by)$ for all $a, b, x, y \in \mathbb{R}$. The threshold duration for **C** is the value $\hat{\tau}(\lambda, c)$ such that $\sigma(\hat{\tau}(\lambda, c)) = c/q$; that is,

$$\hat{\tau}(\lambda, c) \triangleq \frac{1}{\lambda q} \ln\left[\left(1 - \frac{c}{q}\right)^{-1}\right] \qquad \hat{\tau}_q(\cdot, \cdot) < \tau_\lambda(\cdot, \cdot) < 0, \quad \hat{\tau}_c(\cdot, \cdot) > 0 \tag{3.9}$$

Note that we suppress the dependence on q. For $\tau \geq \hat{\tau}(\lambda, c)$, equilibrium values are as in the expressions for generic g_0 after substituting $g_0 = 2$. In this case, the expression of S_0 also simplifies as shown below:

$$V = q\sigma(\tau) - c \tag{3.10}$$

$$S_0^{\text{SIM}} = \sigma(\tau) - \lambda q \tau(1 - \sigma(\tau)) \tag{3.11}$$

Exogenous sequentiality

Proceeding in increasing order of difficulty, we now summarize equilibrium and success rate in the EXO model that we studied in Chapter 2. We will also compute welfare for $g_0 = 2$.

The starting point to compute expected payoffs and welfare expressions is solving the state-contingent success probability evaluated at a funding state $(t, g_t) \equiv (t, g)$,

$$S_{(t,g)} \triangleq \mathbb{E}_{(t,g)}(\mathbb{1}_{\{g_{\tau_+}\}} \leq g_0) \tag{3.12}$$

where $\mathbb{E}_{(t,g)}(\cdot) \equiv \mathbb{E}(\cdot \mid (t, g))$ so that $S_0 \equiv S_{(g_0,0)}$. In Chapter 2 we show that $S_{(t,g)}$ is continuous in time and we derive a recursion that solves equation (3.12) for each $(t, g), g \leq g_0$ given the condition $S_{(\tau,g)} = 1$ for $g \leq 0$.[4] In EXO, the bid conditional success rate satisfies $S_{(t,g)}^{\text{bid}} = S_{(t,g-1)}$ since, as we saw in Section 3.2 when discussing our equilibrium concept, EXO is a Markovian game in which all relevant information for bidders' play is summarized by the funding state. The source of increased difficulty relative to SIM is that now the dynamically evolving funding state matters for an arriving bidder's strategy, in contrast to SIM where bidders strategize based on static campaign

---

[4]Due to continuity this condition would be analogous if stated in terms of τ_+.

attributes.

The bidder expected payoff from playing **C** is now the utility value $U^{\mathbf{C}}_{(t,g)}$ in (3.6) with the twist that S^{bid} is replaced by its state-dependent estimate $S^{\text{bid}}_{(t,g)}$. Again, due to i.i.d. inspection costs and linearity, the equilibrium consists of threshold strategies on cost, but since bidder choice is contingent on the funding state, we will have cutoffs $\hat{c}_{(t,g)}$ for $t \in [0,\tau]$ and $g \in \mathbb{Z}$. Moreover, exogenous sequentiality eliminates the possible coordination failure that instead arises in SIM due to simultaneous actions.

In the homogenous setting, we can classify funding states as *active* if $c \leq \hat{c}_{(t,g)}$ or *frozen* if $c < \hat{c}_{(t,g)}$. We call the second type "frozen" because, if reached, the campaign enters an absorbing state of inactivity. The reason why the campaign becomes permanently inactive is that a frozen campaign enters a vicious circle in which arriving bidders have too little success prospects to find inspecting worthwhile so they never bid and their successors have even worse prospects given that they have less time available for filling the funding gap. Since $S^{\text{bid}}_{(t,g)}$ is decreasing in t, at a given gap g the campaign passes from active to frozen if no bid occurs until remaining duration $\tau - t$ is smaller than a threshold value $\hat{\tau}_g \triangleq \inf \hat{\tau} \in \mathbb{R}_+ : S^{\text{bid}}_{(\tau-\hat{\tau},g)} \geq c/q$. Clearly $\hat{\tau}_g = 0$ for $g \leq 1$, so the equilibrium is fully characterized by solving the non-trivial thresholds $\hat{\tau}_g : S^{\text{bid}}_{(\tau-\hat{\tau}_g,g)} = c/q$ (that is, $U^{\mathbf{C}}_{(\tau-\hat{\tau}_g,g)} = 0$) for all $g > 1$.

Explicit equilibrium and welfare for campaigns with a threshold of two bids

We recap the equilibrium analysis for $g_0 = 2$ and add the welfare calculation. To do so, we only need to analyze bidders' decisions for $g_t = 2$ since bidders are sure that the campaign will succeed for $g_t \leq 1$ so they play **C**. A bidder arriving at $g_t = 2$ instead cannot be sure that a success will occur but does knows that, by bidding, he will convince all his successors, arriving over a period of length $\tau - t$, to also bid. Since success occurs with at last an additional bid from his successors, $S^{\text{bid}}_{(t,2)} = \sigma(\tau - t)$. Using this value, we have that the utility of a bidder playing **C** at $(t, 2)$ is $U^{\mathbf{C}}_{(t,2)} = q\sigma(\tau - t) - c$.

The threshold $\hat{\tau}_2$ solves $\sigma(\hat{\tau}_2) = c/q$ and hence coincides with the minimum required duration $\hat{\tau}(\lambda, c)$, in Eq. (3.9), for an equilibrium with bidding in SIM. In Chapter 2 we denote this threshold by simply $\hat{\tau}_2$. Here instead we let $\hat{\tau}_2 \equiv \hat{\tau}(\lambda, c)$ using a heavier though useful notation since variations in the threshold value (3.9) appear many times in the analysis of ENDO. It follows that, for $\tau < \hat{\tau}(\lambda, c)$, the campaign is stillborn: it starts in a frozen state. For $\tau \geq \hat{\tau}(\lambda, c)$, if $g_t = 2$, bidders play **C** as long as $t \leq \tau - \hat{\tau}(\lambda, c)$ and **A** afterwards; if $g_t \leq 1$ bidders play **C**. The resulting bid rate will decreasing with a discontinuity at $t = \tau - \hat{\tau}(\lambda, c)$ at which the campaign gets frozen if $g_t = 2$.

To compute welfare under the above equilibrium, we already have the success rate expression from Chapter 2 but bidder surplus is missing. We compute it now. In general, bidder surplus in EXO is lower than in SIM, but for the special case of $g_0 = 2$ they are

exactly equal. Now, it is very clear that bidders arriving during $t \in [0, \tau - \hat{\tau}(\lambda, c)]$ achieve the same utility as bidders in SIM since they always inspect and get the good if at least another bidder bids over the whole campaign duration. However, for bidders arriving in the complementary time span $[\tau - \hat{\tau}(\lambda, c), \tau]$ proving the equivalence is not as immediate. Yet, with few algebraic steps, we show that the equivalence still holds.

Lemma 3.1. *The unique* SIM *and* EXO *benchmark equilibria generate the same bidder surplus in threshold two ($g_0 = 2$) campaigns:*

$$V^{\text{EXO}} = V^{\text{SIM}} = q\sigma(\tau) - c \tag{3.13}$$

Proof *in Appendix 3.A.*

To pin-down the success rate, as we already do in Chapter 2, observe that a successful campaign requires at least one bid for $t \in [0, \tau - \hat{\tau}(\lambda, c)]$ and an additional bid in $[\tau - \hat{\tau}(\lambda, c), \tau]$. Given that the Poisson probabilities of no bid by $t = \tau - \hat{\tau}(\lambda, c)$ is $\sigma(\tau - \hat{\tau}(\lambda, c))$, and that of one bid by then and no bid after is

$$\lambda q(\tau - \hat{\tau}(\lambda, c)) \left[1 - \sigma(\tau - \hat{\tau}(\lambda, c))\right] \left[1 - \sigma(\hat{\tau}(\lambda, c))\right] = \lambda q(\tau - \hat{\tau}(\lambda, c)) \left[1 - \sigma(\tau)\right]$$

We can see that the success rate expression is given by

$$S_0^{\text{EXO}} = \sigma(\tau - \hat{\tau}(\lambda, c)) - \lambda q(\tau - \hat{\tau}(\lambda, c))(1 - \sigma(\tau)) \tag{3.14}$$

Equipped with the results developed for benchmark cases, we are ready to solve the simplest campaign under endogenous sequentiality.

3.4   Analysis of crowdfunding with a delay option

We will now introduce the elements required to determine the bidding equilibria and soon move to the analysis of the simplest case in which bidders are homogenous and have inspection cost c. This simple setting already introduces the fundamental elements present in more general variants.

In ENDO bidders are divided into checkers and delayers based on whether they play **C** or **D**. The expected utility a bidder gets by playing **C** is still shown in equation (3.6) but the way $S_{(t,g)}^{\text{bid}}$ is determined is different. The twist is caused by a strictly positive probability of having delayers returning at τ and bidding at the deadline. This changes the end-point condition for $S_{(t,g)}$ in EXO:

$$S_{(\tau,1)} = \sigma(\mu(\mathbf{D})) \qquad \text{and, as usual,} \ \ S_{(\tau,g)} = 0 \ \text{for} \ g > 2 \ \ S_{(\tau,g)} = 1 \ \text{for} \ g \leq 0 \tag{3.15}$$

As anticipated in Section 3.2, $\mu(\mathbf{D})$ will depend on all the stopping times of bids T_b,

$b \in \{0, \dots, g_0 - 2\}$ on the path that the gap takes from state $(0, g_0)$ to $(\tau, 1)$. For this reason now $S_{(t,g)}^{\mathrm{bid}} \neq S_{(t,g-1)}$; that is, the bid-conditional estimate of the success probability made by a bidder in state (t, g) differs from the estimate of the success rate made by an external analyst that observes that the campaign is in state $(t, g-1)$. The reason is that only the bidder knows that the stopping time of his bid is $T_{g_0-g+1} = t$. We will solve $S_{(t,g)}^{\mathrm{bid}}$ explicitly for $g_0 = 2$ in the next subsection.

On the other hand, the expected utility of playing $\mathbf{D}$ depends on the free-riding probability $\phi_{(t,g)}$ and is defined as follows:

$$U_{(t,g)}^{\mathbf{D}} \triangleq \mathbb{E}_{(t,g)}\left(u^{\mathbf{D}}\right) = \phi_{(t,g)}(q - c_t) \quad \phi_{(t,g)} \triangleq \Pr\left(g_\tau \leq 1 \mid (t, g)\right) \tag{3.16}$$

The trade-off between acting upon arrival or postponing depends on $S_{(t,g)}^{\mathrm{bid}}$, which counts as an *assurance probability of success*, and $\phi_{(t,g)}$ which is instead a *free-riding probability*.

3.4.1 Analysis of Endo for a threshold two campaign

In what follows, we will only fully characterize the equilibria of ENDO for $g_0 = 2$, but we will draw general conclusion whenever possible. So, to begin the analysis, we assume $c_t = c \in [0, q]$ and $g_0 = 2$.

The first equilibria that we are going to solve are pure strategy equilibria (p.s.e.) in which bidders play $\mathbf{C}$ if $g_t \leq 1$ and, as long as $g_t = 2$, play $\mathbf{D}$ in a generic set of dates $t \in \mathcal{T}^{\mathbf{D}|2} \subseteq [0, \tau)$ with measure $\delta \triangleq \mu(\mathcal{T}^{\mathbf{D}|2}) \in [0, \tau]$, and play $\mathbf{C}$ in the complementary set $\mathcal{T}^{\mathbf{C}|2} = [0, \tau) \setminus \mathcal{T}^{\mathbf{D}|2}$. Under this class of equilibria, $\mathcal{T}^{\mathbf{D}} = [0, T_1] \cap \mathcal{T}^{\mathbf{D}|2}$ from Eq. (3.4). We call bidders that arrive in $t \in \mathcal{T}^{\mathbf{C}|2}$ *decisives*. This name reflects the fact that they do not need path information in order to decide what to do, so they never delay. Bidders that arrive at $t \in \mathcal{T}^{\mathbf{D}|2}$ are instead potential delayers as they play $\mathbf{D}$ if $g_t = 2$ and $\mathbf{C}$ if $g_t \leq 1$.

To verify that a putative equilibrium of the kind just described exists, we do not have to worry about bidders deviating from $\mathbf{C}$ at $g_t \leq 1$ as they get $U_{(t,1)}^{\mathbf{C}} = U_{(t,1)}^{\mathbf{D}} = q - c \geq 0$, so they do not deviate.[5] The only deviations we need to check carefully are from $\mathbf{C}$ to $\mathbf{D}$ for $t \in \mathcal{T}^{\mathbf{C}|2}$ and vice-versa for $t \in \mathcal{T}^{\mathbf{D}|2}$ in states (t, g) with $g = 2$.

In order to find the that constraints valid equilibria have to satisfy, we determine $S_{(t,2)}^{\mathrm{bid}}$ and $\phi_{(t,2)}$. A bidder playing $\mathbf{C}$ at $(t, 2)$ knows that $T_1 = t$, so only his predecessors will delayed. If follows that the total measure dates at which bidders coincides with the measure of dates over which the bidder's predecessors delayed. In formula, $\mu(\mathbf{D}) = \delta_t^{\mathrm{pre}} \triangleq \mu\left([0, t) \cap \mathcal{T}^{\mathbf{D}|2}\right)$. This value enters $S_{(t,g)}^{\mathrm{bid}}$ and hence $U_{(t,g)}^{\mathbf{C}}$ in the following way:

$$S_{(t,2)}^{\mathrm{bid}} = \sigma\left(\kappa_t + \delta_t^{\mathrm{pre}}\right) \qquad U_{(t,2)}^{\mathbf{C}} = q\sigma\left(\kappa_t + \delta_t^{\mathrm{pre}}\right) - c \tag{3.17}$$

---

[5]We could equally tie-break in favour of $\mathbf{D}$. So doing would not change our results substantively.

with $\kappa_t = \tau - t$ being the remaining duration over which all arriving bidders are checkers.

On the other hand, bidder t manages to free-ride by playing $\mathbf{D}$ if at least one bid is pledged at $g = 2$ by his successors. This bid can come from all of them except delayers arriving during a measure of dates $\delta_t^{\text{post}} \triangleq \mu\left([t,\tau) \cap \mathcal{T}^{\mathbf{D}|2}\right)$.

$$\phi_{(t,2)} = \sigma\left(\kappa_t - \delta_t^{\text{post}}\right) \qquad U_{(t,2)}^{\mathbf{D}} = \sigma\left(\kappa_t - \delta_t^{\text{post}}\right)(q - c) \qquad (3.18)$$

Observing expressions (3.17) and (3.18), we can see that $U_{(t,2)}^{\mathbf{C}}$ exhibits strategic complementarity in $\mathbf{D}$, while $U_{(t,2)}^{\mathbf{D}}$ exhibits strategic substitutability. Most importantly, we can construct the incentive-compatibility constraint that precludes a deviation from $\mathbf{C}$ to $\mathbf{D}$ for $t \in \mathcal{T}^{\mathbf{C}|2}$, denoted $\text{IC}_c^{\mathbf{C},\mathbf{D}}$, by imposing that the difference between the expected utilities in (3.17) and (3.18) is positive.

Interestingly, $\text{IC}_c^{\mathbf{C},\mathbf{D}}$ is stationary. We state this result formally in the next lemma and explain the intuition afterwards.

Lemma 3.2. *The $\text{IC}_c^{C,D}$ condition for willingness to play $\mathbf{C}$ at $(t,2)$ is independent of t:*

$$\sigma\left(\delta\right) \geq c/q, \qquad \delta = \delta_t^{\text{pre}} + \delta_t^{\text{post}} \qquad (\text{IC}_c^{\mathbf{C},\mathbf{D}})$$

Proof *in Appendix 3.A.*

Moves $\mathbf{C}$ and $\mathbf{D}$ yield bidder t the same payoff $q - c$ if at least one bid pledged prior to τ by other arrivals (so that $g_{\tau \leq 1}$) because, either way, he ends up inspecting (at time t if he plays $\mathbf{C}$; at time τ if he plays $\mathbf{D}$). Conversely, no bid from others by τ means that all arriving checkers other than t inspected without bidding but delayers do have the potential to bid still. For this reason $\mathbf{C}$ and $\mathbf{D}$ will result in different realized payoffs: If bidder t plays $\mathbf{D}$, the campaign reaches $g_\tau = 2$ and he earns zero payoff; if he plays $\mathbf{C}$ instead, the campaign reaches $g_\tau = 1$ and the bidder earns q if there is at least another bid from all delayers that arrived during a measure of dates δ. Given that arrivals are independent, by knowing his arrival date bidder t cannot infer any additional information about δ, so also in this case the trade-off between his two possible strategies is independent of his arrival time.

To sum up, the benefit of playing $\mathbf{C}$ relative to $\mathbf{D}$ can be understood in terms of the assurance of success generated by bidding $S_{(\tau,2)}^{\text{bid}} = \sigma\left(\delta\right)$ in a scenario in which all bidding potential of other checkers has been exhausted. What matters ultimately for bidders' decision among $\mathbf{C}$ and $\mathbf{D}$ is the total mass of delaying bidders δ whose contribution potential is at risk of waste. Since $S_{(\tau,2)} = 0$, bidders choose to become decisive only if their pivotality in state $(\tau,2)$, defined as

$$\Delta S_{(\tau,2)} \triangleq S_{(\tau,2)}^{\text{bid}} - S_{(\tau,2)} = S_{(\tau,2)}^{\text{bid}}$$

is sufficiently high relative to the cost-to-quality ratio c/q.

A further important property will help for generalizing the incentive-compatibility condition to heterogenous costs. Specifically, since δ is independent of the inspection cost of any given bidder, we can immediately adapt constraint $(\text{IC}_c^{\mathbf{C},\mathbf{D}})$ so that it applies to a bidder with a generic inspection cost by simply substituting c with c_t. This is valuable and we use it repeatedly in subsections 3.4.2 and 3.4.3.

We have the following putative equilibria: (i) $a_{(t,2)} = \mathbf{D}$ for all t, in which case the campaign necessarily fails. That would imply $\delta = \tau$ and so require $\sigma(\tau) < c/q$ for $\text{IC}_c^{\mathbf{C},\mathbf{D}}$ to never hold. This is possible when $\tau < \hat{\tau}(\lambda, c)$ from Eq. (3.9). Furthermore, also SIM and EXO prescribe inactivity under this condition. (ii) $a_{(t,2)} = \mathbf{C}$ for all t, but that cannot be an equilibrium because an arrival at $(\tau, 2)$ certainly gets a negative payoff from $\mathbf{C}$, which is lower than the zero payoff obtained with $\mathbf{D}$. (iii) on some strict subset of campaign arrival dates, $\mathcal{T}^{\mathbf{D}|2} \subset [0, \tau)$ bidders play $\mathbf{D}$, on the complement play $\mathbf{C}$, and arrivals are indifferent i.e. both $\text{IC}_c^{\mathbf{C},\mathbf{D}}$ and $\text{IC}_c^{\mathbf{D},\mathbf{C}}$ hold, or simply $U_{(t,2)}^{\mathbf{C}} = U_{(t,2)}^{\mathbf{C}}$ for all t.

Since (iii) is the only possibility and expected utilities equate at $\delta = \hat{\tau}(\lambda, c)$, we have characterized all pure strategy equilibria.[6]

For mixed strategy equilibria (m.s.e.), note that if players mix $\rho_t \in (0, 1)$ at t on $\mathbf{C}$ and $1 - \rho_t$ on $\mathbf{D}$, then the only equilibrium restriction is $\int_0^\tau (1 - \rho_t) \, \mathrm{d}t = \hat{\tau}(\lambda, c)$, giving a rich continuum of possibilities. The unique stationary equilibrium is the m.s.e. where $\rho_t = \rho \; \forall t$,

$$\rho = (\tau - \hat{\tau}(\lambda, c))/\tau \tag{3.19}$$

These statements reflect the required incentive-compatibility condition under mixing and complete the proof of the next proposition on equilibrium analysis.

Proposition 3.1. *If $\tau < \hat{\tau}(\lambda, c)$, ENDO (like EXO and SIM) has only trivial equilibria in which there is never any bidding. If instead $\tau \geq \hat{\tau}(\lambda, c)$, arrivals always play $\mathbf{C}$ if $g_t \leq 1$, and we characterize the equilibria by describing the possible strategies at states $(t, 2)$:*

(a) Every p.s.e. is described by a subset $\mathcal{T}^{\mathbf{D}|2} \subset [0, \tau)$, on which $a_{(t,2)} = \mathbf{D}$ for all $t \in \mathcal{T}^{\mathbf{D}|2}$ and $a_{(t,2)} = \mathbf{C}$ otherwise. The subset $\mathcal{T}^{\mathbf{D}|2}$ must have measure $\hat{\tau}(\lambda, c)$.

(b) In any m.s.e., the set of dates at which $\mathbf{D}$ would be played if the gap stayed at $g = 2$ throughout has expected measure $\hat{\tau}(\lambda, c)$.

(c) There is a unique equilibrium in stationary strategies. In this equilibrium, any arrival facing $g_t = 2$ mixes with probability weight $\rho = \frac{\tau - \hat{\tau}(\lambda, c)}{\tau}$ on $\mathbf{C}$ and $1 - \rho$ on $\mathbf{D}$.

We draw attention to the specific p.s.e. that prescribes $\mathcal{T}^{\mathbf{D}|2} = [\tau - \hat{\tau}(\lambda, c), \tau]$. This equilibrium is equivalent to the unique equilibrium in EXO because while that equilibrium has bidders playing $\mathbf{A}$ instead of $\mathbf{D}$ for states $(t, 2)$ with t on that same interval, $\mathbf{D}$ is

---

[6]We can safely ignore strategy $\mathbf{D\bar{2}}$ as it does not survive the 1S-1W refinement. First any strategy where $\mathbf{A}$ is played for $g_t \leq 1$ is strictly dominated by others that prescribe choosing $\mathbf{C}$. So we are left with only $\mathbf{C}$ played for $g_t \leq 1$. Now, if $\mathbf{D}$ is not played in equilibrium, strategies $\mathbf{C}$ and $\mathbf{D\bar{2}}$ yields the same payoff. Otherwise, $\mathbf{C}$ payoff is strictly higher due to the encouragement effect of bidding on delayers. So $\mathbf{D\bar{2}}$ is eliminated in the second step of the refinement criterion.

equivalent to **A** given that there is no chance of a bid bringing the gap to 1 or lower and activating the delayers.

Corollary 3.1. *The p.s.e. with $\mathcal{T}^{\mathbf{D}|2} = [\tau - \hat{\tau}(\lambda, c), \tau)$ is profile-equivalent to the unique equilibrium of* Exo*.*

The previous proposition and its corollary characterized the equilibrium set of the homogenous case for $g_0 = 2$. It is important to understand however that this equilibrium structure is very specific to the case of campaigns with a goal of two bids. E.g. for $g_0 = 3$ a bid at $g_t = g_0$ does not motivate all bidder successors and current delayers to inspect, so the intuition on pivotality on the terminal campaign does not hold any more. Moreover, the number of potential delayers that can wait at $g = 2$ changes according to the stopping time of the first bid.

Welfare calculations

We now compute success rate and welfare achieved in the equilibria of Proposition 3.1 and compare them with their values attained in Sim and Exo. We find that sequentiality of actions reduces the success rate as it allows bidders to free-ride, but, for $g_0 = 2$, it is irrelevant for bidder surplus. We start by proving the latter result.

The bidder surplus can be determined based on the observation that any bidder with $t \in \mathcal{T}^{\mathbf{D}|2}$ ends up inspecting at t or τ if at least one bid is pledged by checkers. Hence the stopping time of the first bid has no effect on his utility. Since equilibrium requires indifference between **C** and **D** strategies (which holds while $g_t = 2$ *and* while $g_t = 1$), per-bidder surplus is just this expected payoff of action **D** computed from date 0. Given that all **D** bidders inspect if **C** bidders pledge at least one bid during the whole duration, with probability $\sigma(\rho\tau) = \sigma(\tau - \hat{\tau}(\lambda, c))$, all equilibria with activity presented in Proposition 3.1 produce the same bidder surplus of

$$V^{\text{Endo}} = \lambda\tau(q - c)\sigma(\tau - \hat{\tau}(\lambda, c)), \quad \tau - \hat{\tau}(\lambda, c) = \rho\tau \tag{3.20}$$

Here we can quickly grasp that, in the homogeneous case, assumptions on bidders' timing of actions are totally neutral on their expected surplus for $g_0 = 2$ campaigns: The per-bidder surplus of delayers is in fact equivalent to the one of bidders in Exo that arrive after the critical date $\tau - \hat{\tau}(\lambda, c)$: both types inspect only if at least one bid is pledged during a measure $\tau - \hat{\tau}(\lambda, c)$ of dates, and in that case they are sure that the campaign will succeed after they bid. We can conclude that Sim, Exo and Endo produce the same bidder surplus.

The same indifference result is not valid when discussing the success rate. For $\tau \geq \hat{\tau}(\lambda, c)$, Sim has S^{bid} exceeding c/q and therefore *all* bidders are willing to play **C**. In Endo instead, **D** bidders may never inspect. This reduces the success rate. The

precise value of S^{ENDO} is the same for all p.s.e. and m.s.e. with activity described by Proposition 3.1, which follows by subtracting one from the probability of no bid pledged by **C** bidders, plus the probability of exactly one bid from **C** bidders but no bid from **D** returners.

$$S_0^{\text{ENDO}} = \sigma\,(\tau - \delta) + \lambda q(\tau - \delta)(1 - \sigma\,(\tau)) \tag{3.21}$$

Given that the measure of dates in which delays occur coincides with the one of Exo bidders arriving beyond the critical duration threshold, i.e. $\delta = \hat{\tau}(\lambda, c)$, we have demonstrated that $S_0^{\text{ENDO}} = S_0^{\text{EXO}}$ from Eq. (3.14). We can draw the following conclusion:

Combining Eq. (3.20) and Eq. (3.21), we can see that sequentiality weakly reduces welfare in the homogenous case.

Proposition 3.2. *In the homogenous case, for $g_0 = 2$ campaigns, we have $V = S_0 = W = 0$ in all settings for $\tau < \hat{\tau}(\lambda, c)$. In the opposite case instead,*

(i) $V^{\text{ENDO}} = V^{\text{EXO}} = V^{\text{SIM}}$;

(ii) $S^{\text{ENDO}} = S^{\text{EXO}} < S^{\text{SIM}}$;

(iii) $W^{\text{ENDO}} = W^{\text{EXO}} \leq W^{\text{SIM}}$, with the last inequality strict for $R > 0$.

This result already anticipates part of the optimal design analysis in Section 3.5 since clearly the simultaneous (no-information) design is optimal under homogeneous inspection costs. As we will see next however, the trade-off between sequential and simultaneous play becomes more subtle when we introduce cost heterogeneity. Indeed, when the crowdfunding campaign is hard, we find that endogeneity raises welfare.

3.4.2 Binary inspection costs

Now we analyze how heterogeneity affects the results of the previous section. To start with, the binary inspection cost setting provides the most elementary form of heterogeneity. Then we show that some features of the binary setting still hold with a continuum of inspection costs.

We assume that the inspection cost CDF has support $\{c_L, c_H\} : q \geq c_H > c_L \geq 0$ and assigns a probability z on c_L. We already solved $c_L = c_H$ in Section 3.4.1, so we can ignore this case. The binary setting requires two distinct incentive-compatibility conditions; one for each cost-type bidder. These follow from Eq. ($\text{IC}_c^{\mathbf{C},\mathbf{D}}$) after substituting c with c_j, $j \in \{L, H\}$. The pivotality of a bidder on the campaign in state $(\tau, 2)$ is still computed based on the total *expected* mass of delayers, which now can be decomposed into measures of arrival dates of each cost group: $\delta = z\delta_L + (1 - z)\delta_H$. So, for each $j \in \{L, H\}$, the new incentive-compatibility constraints is

$$\sigma\,(z\delta_L + (1 - z)\delta_H) \geq c_j/q, \tag{$\text{IC}_{c_j}^{\mathbf{C},\mathbf{D}}$}$$

We will refer to the constraint of a specific cost type as $\text{IC}_{c_L}^{\mathbf{C},\mathbf{D}}$ and $\text{IC}_{c_H}^{\mathbf{C},\mathbf{D}}$.

As before, p.s.e. specify two sets of dates $\mathcal{T}_L^{\mathbf{D}|2}$, $\mathcal{T}_H^{\mathbf{D}|2}$, with $\delta_L = |\mathcal{T}_L^{\mathbf{D}|2}|$ and $\delta_H = |\mathcal{T}_H^{\mathbf{D}|2}|$, at which L and H types delay if $g_t = 2$. These dates must be incentive-compatible for both cost-types simultaneously. M.s.e. specify mixing weights and in those cases we define δ_L and δ_H as the corresponding expected measures of $g=2$-contingent delay times.

The equilibrium set again features a large multiplicity as we will see soon, but to be concise, we present only equilibria in stationary strategies. These involve probability weights on $\mathbf{C}$ given by $\{\rho_L, \rho_H\}$ with $\rho_j \in [0,1]$ for each $j \in \{L, H\}$. For equilibria in stationary strategies, a type j only plays a pure strategy if $\rho_j = 0$ or 1. We only briefly point out the possibilities for equilibrium multiplicity alongside this since the extended set is analogous to that under homogeneity.

Before moving forward, we point out that condition $(\text{IC}_{c_j}^{\mathbf{C},\mathbf{D}})$ shows a clear monotonicity of incentives with respect to inspection costs: since the same δ appears in both constraints but $c_L < c_H$, we have that if some H-type is decisive, then all L-types are decisive, while if not all L-types are decisive, then none of the H-type is decisive. This is easily proved by the fact that, if $\text{IC}_{c_H}^{\mathbf{C},\mathbf{D}}$ is binding ($\rho_H > 0$), $\text{IC}_{c_L}^{\mathbf{C},\mathbf{D}}$ is slack ($\rho_L = 1$); conversely, if $\text{IC}_{c_L}^{\mathbf{C},\mathbf{D}}$ is binding ($\rho_L > 0$), $\text{IC}_{c_H}^{\mathbf{C},\mathbf{D}}$ breaks ($\rho_H = 0$). This observation proves the next lemma which will guide our analysis of equilibria in the binary case.

Lemma 3.3. $\rho_H > 0$ *implies* $\rho_L = 1$; $\rho_L < 1$ *implies* $\rho_H = 0$

To find the inactivity equilibrium in the binary setting, it is sufficient to check that $\rho_L = 0$ since that implies $\rho_H = 0$. This happens by breaking $\text{IC}_{c_L}^{\mathbf{C},\mathbf{D}}$ which occurs for any $\tau < \hat{\tau}(\lambda, c_L)$. So the last condition ensures inactivity. Equilibria with inactivity do not exist for $c_L = 0$ ($\hat{\tau}(\lambda, 0) = 0 \geq \tau$).

The second type of equilibrium in stationary strategies has $\rho_L \in (0,1)$ and still $\rho_H = 0$ (L-types mix and H-types always delay). $\text{IC}_{c_L}^{\mathbf{C},\mathbf{D}}$ is binding (as well as $\text{IC}_{c_L}^{\mathbf{D},\mathbf{C}}$) but $\text{IC}_{c_H}^{\mathbf{C},\mathbf{D}}$ fails. It follows that $\delta = z\delta_L + (1-z)\tau = \hat{\tau}(\lambda, c_L)$. We solve for the value δ_L that is consistent with this equilibrium and find

$$\delta_L = \hat{\tau}(\lambda z, c_L) - \tau \frac{1-z}{z}, \ \rho_L = \frac{\tau - \hat{\tau}(\lambda z, c_L)}{\tau} + \frac{1-z}{z} \in (0,1) \qquad (3.22)$$

Formula (3.22) shows that ρ_L is decreasing in both c_L and $\lambda_L \triangleq \lambda z$ and increasing in the (H-to-L-types) participation ratio $(1-z)/z$. Intuitively, higher c_L makes free-riding more attractive on the savings side since L-types lose more when sinking their cost in vain on a campaign that fails, while λ_L raises assurance by raising the expected number of delayers for a given ρ_L. In the other direction, a large $(1-z)/z$ deters free-riding because it makes bidders more pivotal as their bid convinces not only all L-types who delay but also all H-type types to eventually inspect.

As incentive constraints are independent of time, this equilibrium in stationary strategies lies within a broad continuum of other valid m.s.e. and p.s.e. with $\mathbb{E}(\mathcal{T}_L^{\mathbf{D}|2}) = \delta_L$.

The slight difference with respect with the p.s.e. in Proposition 3.1 is that δ_L varies (decreases) continuously with τ.

For $\tau \geq \overline{\tau}_L$, where

$$\overline{\tau}_L \triangleq \hat{\tau}(\lambda(1-z), c_L) \tag{3.23}$$

the encouragement effect of a bid on the H-types is so strong that $\text{IC}_{c_L}^{\mathbf{C,D}}$ holds even at $\delta_L = 0$. Therefore, in the range $\tau \geq \overline{\tau}_L$ we have $\rho_L = 1$ and $\rho_H \geq 0$ (H-types mix and all L-type play $\mathbf{C}$); which means $\delta = (1-z)\delta_H$.

Having all the L-types decisive is still not enough to make H-types confident to inspect. It follows that we have to further distinguish two sub-cases for this parameter region. We can have the following cases. First, $\rho_H = 0$ and $\delta = (1-z)\tau$. This requires $\tau < \hat{\tau}(\lambda(1-z), c_H)$, so we have a unique stationary equilibrium $\{\rho_L, \rho_H\} = \{1, 0\}$ for $\tau \in [\hat{\tau}(\lambda, c_L), \hat{\tau}(\lambda(1-z), c_H))$. Second, $\rho_H > 0$ and $\delta = (1-z)\delta_H$, with $\delta_H < \tau$. We have this equilibrium for $\tau \geq \hat{\tau}(\lambda(1-z), c_H)$ so that δ_H is determined according to the binding constraint $\text{IC}_{c_H}^{\mathbf{C,D}}$:

$$\delta_H = \hat{\tau}(\lambda(1-z), c_H), \quad \rho_H = \frac{\tau - \hat{\tau}(\lambda(1-z), c_H)}{\tau} \in (0, 1) \tag{3.24}$$

Of course this equilibrium is again part of a continuum of valid predictions that satisfy $\mathbb{E}\left(\mathcal{T}_H^{\mathbf{D}|2}\right) = \delta_H$.

Some remarks are in order at this point. First, H-types free-riding incentives depend only on the arrival rate of their own type $\lambda_H = \lambda(1-z)$ because all L-type choose to be decisive in this parameter range. Also, it is curious to see that $(1-z)$ appears in both cutoffs (3.24) and (3.23), yet it does so for two clear and different reasons. In (3.24), it comes from an H-type *intra-group* encouragement effect on other H-types. In (3.23) instead, that term appears as a result of *infra-group* encouragement that a bid from an L-type provokes on H-types who arrive with probability $1 - F(c_L) = 1 - z$. This final clarification together with the earlier explanation concludes the analysis of equilibria in stationary strategies in the binary setting and proves the following summary result.

Proposition 3.3. *The mixing weights (ρ_L, ρ_H) for a stationary equilibrium take the following values.*

$\{\rho_L, \rho_H\} = \{0, 0\}$ *for* $\tau < \hat{\tau}(\lambda, c_L)$;
$\rho_L \in (0, 1), \rho_H = 0$, *with ρ_L in Eq. (3.22), for* $\tau \in [\hat{\tau}(\lambda, c_L), \hat{\tau}(\lambda(1-z), c_L))$;
$\{\rho_L, \rho_H\} = \{1, 0\}$ *for* $\tau \in [\hat{\tau}(\lambda(1-z), c_L), \hat{\tau}(\lambda(1-z), c_H)]$;
$\rho_L = 1, \rho_H \in (0, 1)$, *with ρ_H in Eq. (3.24), for* $\tau \in (\hat{\tau}(\lambda(1-z), c_H), +\infty)$.

The computation of welfare under the equilibria in Proposition 3.3 is straightforward. Per-bidder surplus can we expressed as the average surplus across cost types

$$V = zV_L + (1-z)V_H \tag{3.25}$$

Now, obviously the inactivity equilibrium produces zero surplus (and welfare), so we focus on more interesting cases. In all other equilibria, *L*-type are either indifferent or better-off by playing **C**. So their surplus is simply

$$V_L = q\sigma\left(\tau\right) - c_L \tag{3.26}$$

By contrast, *H*-type are either indifferent or strictly prefer **D**, so

$$V_H = (q - c_H)\sigma\left(\tau - \delta\right), \delta = \begin{cases} \hat{\tau}(\lambda, c_L) & \text{for} \tau \in [\hat{\tau}(\lambda, c_L), \hat{\tau}(\lambda(1-z), c_L)) \\ \tau(1-z) & \text{for} \tau \in \tau \in [\hat{\tau}(\lambda(1-z), c_L), \hat{\tau}(\lambda(1-z), c_H)] \\ \hat{\tau}(\lambda, c_H) & \text{for} \tau \in (\hat{\tau}(\lambda(1-z), c_H), +\infty) \end{cases} \tag{3.27}$$

The success rate can still be generically stated as

$$S_0 = \sigma\left(\tau - \delta\right) + \lambda q(\tau - \delta)(1 - \sigma\left(\tau\right)) \tag{3.28}$$

where δ is again determined by the cases displayed in Eq. (3.27).

3.4.3   Continuum of inspection costs

A broader form of heterogeneity than binary is to have inspection costs drawn from a continuous distribution function F over $c \in [0, q]$. Given that incentives are monotone with respect to inspection costs, equilibria are again in threshold strategies. There is a threshold $\hat{c}$, such that bidders play **C** for $c \leq \hat{c}$ and **D** if not. It follows that, in expectation, $\delta = \tau[1 - F(\hat{c})]$, so that the incentive-compatibility condition for **C** of a type c_t is $\sigma\left(\tau(1 - F(\hat{c}))\right) \geq c_t/q$, while that for **D** has the inequality reversed.

The threshold type is indifferent between **C** and **D**, so both constraints bind for this types. The cutoff costs $\hat{c}$ that can form an equilibrium are defined implicitly as the solution to the following equation.

$$\sigma\left(\tau(1 - F(\hat{c}))\right) = \hat{c}/q \tag{IC$_{\hat{c}}$}$$

Finding equilibria boils to determine $\hat{c}$ values for which equation (IC$_{\hat{c}}$) holds. With the proof of the next lemma, we show that the technical conditions for a unique solution and interior solution $\hat{c} \in (0, q)$ to Eq. (IC$_{\hat{c}}$) hold. So we have a unique interior and stationary equilibrium.

Proposition 3.4. *With a continuous CDF $F(\cdot)$ over $c_t \in [0, q]$, for $g_0 = 2$ campaigns, there exists a unique, interior and stationary p.s.e. in threshold strategies $a_{(t,g)} = \boldsymbol{C}$ if $c_t \leq \hat{c}$ and $a_{(t,g)} = \boldsymbol{D}$ otherwise. The cutoff $\hat{c} \in (0, q)$ is the unique solution to Eq. (IC$_{\hat{c}}$).*

Proof *in Appendix 3.A.*

The expressions of payoffs, welfare and success rate in Section 3.4.1 under the stationary m.s.e. hold also here after applying the substitution $\rho = F(\hat{c})$ so that $S^{\text{bid}}_{(t,2)} = \sigma\left(\tau - tF(\hat{c})\right)$ and $\phi_{(t,2)} = \sigma\left((\tau - t)F(\hat{c})\right)$. It follows that the success rate is given by

$$S_0 = \sigma\left(F(\hat{c})\tau\right) - \lambda q F(\hat{c})\tau(1 - \sigma\left(\tau\right)) \tag{3.29}$$

3.5 Optimal design

In this section, we compare the baseline ENDO model against the two benchmarks. Recall that a platform design that provides no information is equivalent to SIM, so it is meaningful to compare ENDO against SIM with bidders arriving over time according to exactly the same Poisson process. The relevance of ENDO versus EXO as models of crowdfunding is partly determined by bidder time constraints, preferences and decision-making capabilities. Nonetheless, "Remind-me" buttons facilitate the delay strategy, giving platform design a clear role. It is therefore meaningful to treat the comparison as a design question, with the proviso that both ENDO and EXO are extremes and the comparison exaggerates the platform's degree of control. Later, we explain how to generalize the models by allowing some bidders to choose to delay while others act on arrival or never.

We begin with a hypothetical optimal design in which we assume bidders are willing to play **C** if that benefits the social interest. We call this the "constrained first-best" setting since we impose the constraint that bidders are only able to communicate via the aggregate bid level that is made transparent on the crowdfunding platform.

3.5.1 Constrained first-best

Let $*$ denote the constrained optimum with no free-riding. We solve W^* where a social planner can oblige bidders to play any strategies the planner wishes, but has to respect their costs of information and the communication constraints defined by the crowdfunding scenario. The information constraints are exactly as in the baseline model: a bidder t only knows his valuation if he inspects by paying his cost c_t. The communication constraint restricts what information the planner can use in order to decide what strategies to have bidders follow. A planner would ideally select the socially optimal bidder actions based on how many bidders arrive and the value of their inspection costs, but that would give the planner an exaggerate degree of control over bidder actions since those bidder characteristics are private and hardly disclosed. More useful is the benchmark in which the planner dictates to bidders what strategy to follow without affecting any interactive

communication beyond the interaction implemented by the crowdfunding platform.

In particular, we look for the (socially) optimal number of delaying bidders when the planner can impose to a first group of bidders arriving during $t \in (0, \hat{t}^*)$, for some $t^* \in (0, \tau]$, to become decisive and play **C** and to a second group arriving during the remaining time $t \in [\hat{t}^*, \tau]$ to play **D**. In this way all bidders within each action group moves simultaneously (even though privately they may not be willing to do it). We assume $\tau \geq \hat{\tau}(\lambda, c)$ to study the non-trivial case. We let

$$V_1 = q\sigma(\tau) - c \qquad V_2 = (q - c)\sigma(\tau_1) \qquad V = \frac{\tau_1}{\tau_1 + \tau_2}V_1 + \frac{\tau_2}{\tau_1 + \tau_2}V_2 \qquad (3.30)$$

denote the per-bidder surplus of bidders that belong to the group of decisives and delayers and the average per-bidder surplus, with $\tau_1 = \mu((0, t^*))$ and $\tau_2 = \mu([t^*, \tau])$. The optimal constrained design are values τ_1, τ_2 that maximize V in (3.30) subject to $\tau_1 + \tau_2 = \tau$. Since arrivals of both bidder groups cover the whole, we have one degree of freedom only to maximize the average surplus in Eq. (3.30). The constrained first-best solves the following problem.

$$V^* = \max_{\tau_1} \tau_1 \left(q\sigma(\tau) - c\right) + (\tau - \tau_1)(q - c)\sigma(\tau_1) \qquad (3.31)$$

It is straightforward to verify that the maximization problem equation (3.31) is regular and admits a unique solution using the first- and second-order conditions. Most importantly, we prove that neither SIM nor ENDO are optimal designs:SIM has no delaying bidders, which are too few for satisfying optimality; ENDO on the other hand is distorted towards the opposite extreme of inducing too many bidders to delay. The source of the ENDO distortion is clear: the planner takes into account the positive externality of decisives on delayers and therefore raises the expected number of decisives. In the ENDO equilibrium, all bidders are indifferent between delay and decisively inspecting. Marginally raising the measure of bidders who are decisive has a first-order benefit on those who delay and at most a second-order cost on those who become decisive; in fact, the cost on decisives relative to the equilibrium is zero. It follows that the optimal degree of delay is less than the equilibrium value in ENDO.

The SIM distortion result is also straightforward. From a setting in which no bidder delays, increasing the measure of delay imposes no cost on any existing bidder since those who inspect are unaffected: they know that delayers will inspect at the deadline if they bid and they only value other bidders' inspecting in the scenario where they want to bid and therefore care about success. In conclusion, the constrained optimum has less decisiveness than does SIM. That is, some delay is optimal, but less than in ENDO where free-riding lowers the success rate, wasting some of the potential for social gains.

Proposition 3.5. SIM *and* ENDO *respectively imply insufficient and excessive delay rel-*

ative to the optimal communication-constrained design.

Proof *in Appendix 3.A.*

3.5.2 Comparing Endo vs. Sim and Exo given heterogeneity

We saw that in the baseline case of homogenous inspection costs sequentiality is harmful to welfare. Nevertheless, in this section, we will prove that heterogeneity gives an advantage to endogenous sequentiality relative to exogenous sequentiality or to simultaneous play, which as a result can provide higher welfare. We will provide some examples where we compare ENDO to EXO and SIM. Before doing so, we present the simple equilibria of SIM with binary costs for $g_0 = 2$. EXO is already solved in Chapter 2 and we do not need explicit results on that for the comparisons that will follows. In order to compare design options, we let explicitly $\hat{c}^{\text{SIM}}$ and $\hat{c}^{\text{ENDO}}$ denote the threshold costs when needed.

Sim with heterogenous costs

Equilibria of SIM for a $g_0 = 2$ campaign are very easy to find. If all bidders participate, each gets $U^{\mathbf{C}} = q\sigma(\tau) - c_t$. We analyze the binary case first: $c_t \in \{c_L, c_H\}$. For $\tau \geq \hat{\tau}(\lambda, c_H)$ all types are happy to play $\mathbf{C}$. So for this range of values, S_0 is once again given by Eq. (3.11). Surplus follows by substituting $E[c_t]$ to c in expression Eq. (3.10). So,

$$V^{\text{SIM}} = q\sigma(\tau) - (zc_L + (1-z)c_H) \tag{3.32}$$

For $\tau \in [\hat{\tau}(\lambda z, c_L), \hat{\tau}(\lambda, c_H))$, the H-types cannot obtain a positive payoff when all bidders inspect together. Still, L-types achieve a positive utility if all of them inspect simultaneously. So H-types play $\mathbf{A}$ and L-types play $\mathbf{C}$ for this parameter range. Since only L-type bid, their average rate of bidding is only λz, so per-bidder surplus and success rate are

$$V^{\text{SIM}} = z(q\sigma(z\tau) - c_L) \qquad S_0^{\text{SIM}} = \sigma(z\tau) + \lambda qz\tau(1 - \sigma(z\tau)) \tag{3.33}$$

Finally For $\tau < \hat{\tau}(\lambda z, c_L)$, not even the L-type gain positive expected utility by participating, so the crowdfunding campaign remains inactive. In this case, of course, both bidder surplus and success rate are zero.

SIM with a continuum of costs $c_t \in [0, q]$ distributed according to a generic $F(\cdot)$ has equilibria in threshold strategies of the kind mentioned in Section 3.3 in which bidders choose $\mathbf{C}$ for $c \leq \hat{c}$ and $\mathbf{A}$ for $c > \hat{c}$. The threshold $\hat{c}$ makes bidders indifferent between the two strategies. That is,

$$\sigma(F(\hat{c})\tau) = \hat{c}/q \tag{IC$^{\text{SIM}}$}$$

In principle, (IC$^{\text{SIM}}$) can have multiple solution among which $\hat{c} = 0$. In some cases, $\hat{c} = 0$ is the unique solution and in that case the campaign is inactive. When instead there are also others, Pareto optimality selects the highest. Despite all these details and the fact that moreover we cannot solve explicitly for $\hat{c}$, all we need to compare SIM with EXO is just (IC$^{\text{SIM}}$), so we do not dig further into understanding these details.

Given a value $\hat{c}$, the success rate is the probability of at least two simultaneous bids from $\lambda F(\hat{c})\tau$ bidders on average, each bidding with probability q.

$$S_0^{\text{SIM}} = \sigma\left(F(\hat{c})\tau\right) - \lambda q F(\hat{c})\tau \left[1 - \sigma\left(F(\hat{c})\tau\right)\right] \tag{3.34}$$

Welfare comparisons with heterogenous inspection costs

We start with the binary setting. We can quickly prove that $V^{\text{ENDO}} \geq V^{\text{SIM}}$. First, ENDO has inactivity in a smaller parameter range $\tau < \hat{\tau}(\lambda, c_L)$ since $\hat{\tau}(\lambda, c_L) > \hat{\tau}(\lambda z, c_L)$. For $\tau < \hat{\tau}(\lambda, c_L)$, both models give zero welfare. However $\tau \in [\hat{\tau}(\lambda, c_L), \hat{\tau}(\lambda z, c_L))$ ENDO has bidder participation so produces positive welfare; SIM instead prescribes inactivity.

Next, for $\tau \in [\hat{\tau}(\lambda z, c_L), \hat{\tau}(\lambda, c_H))$, SIM has H-types not participating whereas ENDO has both L- and H- types participating. Moreover, L-types utility in ENDO is strictly higher than in SIM since when bidding they encourage also the H-type types whereas this does not happen in SIM.

Finally, for $\tau \geq \hat{\tau}(\lambda, c_H)$ both types inspect in SIM. In this case bidder surplus in ENDO is again weakly above its value in SIM because H-types can obtain an information rent by delaying, which rises their surplus relative to SIM. In all the cases that we analyzed, we have $V^{\text{ENDO}} \geq V^{\text{SIM}}$.

Now we compare the success probabilities implied by the alternative design options. As we already know, simultaneity raises the success rate with respect to sequentiality if all bidders bid at the same time. However, when bidders that do not inspect simultaneously in SIM instead do inspect after delaying in ENDO, endogenous sequentiality can even raise the success rate.

To be specific, for $\tau \geq \hat{\tau}(\lambda, c_H)$, SIM has all bidders inspecting, hence achieves the maximal success rate S_0^{SIM} in expression Eq. (3.11). For $\tau \in [\hat{\tau}(\lambda z, c_L), \hat{\tau}(\lambda, c_H))$, the H-types play **A** in SIM but play either **D** or **C** in ENDO, so as long as all the L-type play **C** (as in SIM), ENDO gives a strictly higher success probability. This requires $\hat{\tau}(\lambda(1-z), c_L) \geq \hat{\tau}(\lambda z, c_L)$ and hence $z < 1/2$. In other words, this condition ensures that the probability of H-type arrivals $1 - z$ is high enough to give L-types sufficient pivotality. The case of $\tau < \hat{\tau}(\lambda z, c_L)$ is nice and extreme since ENDO can give a chance of success to the campaign while SIM cannot. We have proved.

Proposition 3.6. *For binary inspection costs, $V = S_0 = W = 0$ in all three models if $\tau < \hat{\tau}(\lambda, c_L)$. In the opposite case,*

(i) $V^{\text{ENDO}} \geq V^{\text{SIM}}$;

(ii) For $\tau \geq \hat{\tau}(\lambda, c_H)$, $S_0^{\text{ENDO}} < S_0^{\text{SIM}}$; for $\tau < \hat{\tau}(\lambda, c_H)$ and $z < 1/2$, $S_0^{\text{ENDO}} > S_0^{\text{SIM}}$;

(iii) For $\tau \geq \hat{\tau}(\lambda, c_H)$ and R sufficiently small, $W^{\text{ENDO}} \geq W^{\text{SIM}}$; for $\tau < \hat{\tau}(\lambda, c_H)$, and $z < 1/2$, $W^{\text{ENDO}} \geq W^{\text{SIM}}$, with strict inequality for any positive R.

Proof *in Appendix 3.A*.

Notice that EXO also beats SIM in case $\tau < \hat{\tau}(\lambda, c_H)$ and $z < 1/2$ since it gives a chance to have H-types inspecting. However, it is also important to see that ENDO strictly beats EXO in that case because it does not waste the H-types who arrive prior to T in EXO (in which T is exponentially distributed with parameter λz). In cases other than that just presented, the comparison of EXO to the other design options is more involved since it can go in either direction, depending on parameters.

The main design insights for binary types hold as well with a continuum of inspection costs. Going into the details, since **C** yields the same surplus in both SIM and ENDO, we have that types $c_t \leq \hat{c}^{\text{ENDO}}$ get always a surplus of $q\sigma(\tau) - c_t$ in ENDO while in SIM those with $c_t \in (\hat{c}^{\text{SIM}}, \hat{c}^{\text{ENDO}}]$ get zero. Therefore, if $\hat{c}^{\text{ENDO}} > \hat{c}^{\text{SIM}}$, ENDO raises surplus of those types. On the other hand, all types $c_t \geq \hat{c}^{\text{ENDO}}$ instead would obtain the same surplus they obtain in SIM if they were playing **C** in ENDO, but they actually play **D** which gives them a higher surplus (type $c_t = \hat{c}^{\text{ENDO}}$ is indifferent among the two design options; types $c_t > \hat{c}^{\text{ENDO}}$ strictly prefer ENDO). Hence, in ENDO all bidders are at least as well-off as in SIM.

We also find that small enough τ is sufficient for $S_0^{\text{ENDO}} > S_0^{\text{SIM}}$. With a continuous $F(\cdot)$, the last term of expressions Eq. (3.34) and Eq. (3.29) show that $S_0^{\text{ENDO}} \geq S_0^{\text{SIM}}$ if $\hat{c}^{\text{ENDO}} \geq \hat{c}^{\text{SIM}}$. This means that the number of bidders who avoid bidding under SIM is higher than the number of delayers in ENDO. Bearing in mind that the equilibrium thresholds are functions of the fundamentals, for the next few steps only, we let $\hat{c}^{\text{ENDO}} \equiv \hat{c}^{\text{ENDO}}(\lambda q\tau)$; $\hat{c}^{\text{SIM}} \equiv \hat{c}^{\text{SIM}}(\lambda q\tau)$. Equilibrium conditions (IC$^{\text{SIM}}$) and (IC$_{\hat{c}}$) show that $\hat{c}^{\text{ENDO}}(\lambda q\tau) \geq \hat{c}^{\text{SIM}}(\lambda q\tau)$ if and only if $1 - F(\hat{c}^{\text{ENDO}}(\lambda q\tau)) \geq F(\hat{c}^{\text{SIM}}(\lambda q\tau))$ or

$$F(\hat{c}^{\text{ENDO}}(\lambda q\tau)) + F(\hat{c}^{\text{SIM}}(\lambda q\tau)) \leq 1 \tag{3.35}$$

As the two cutoffs are both increasing τ and (λ, q), both equal to 0 when $\lambda q\tau = 0$ and strictly positive otherwise and F is continuous, inequality (3.35) holds when $\lambda q\tau$ are sufficiently small, hence satisfied for τ small. Fig. 3.1 illustrates the effect of parameter shifts on the equilibrium cutoffs.

The comparison is not immediate in case $\hat{c}^{\text{SIM}} < \hat{c}^{\text{SIM}}$. Nevertheless, we already made the point that ENDO helps raising success rate and welfare if $\lambda q\tau$ is sufficiently low, whereas SIM does better in the opposite case.

The analysis of this section demonstrated a valid principle for threshold two campaigns: with either type of heterogeneity, a delay option design always outperformed the
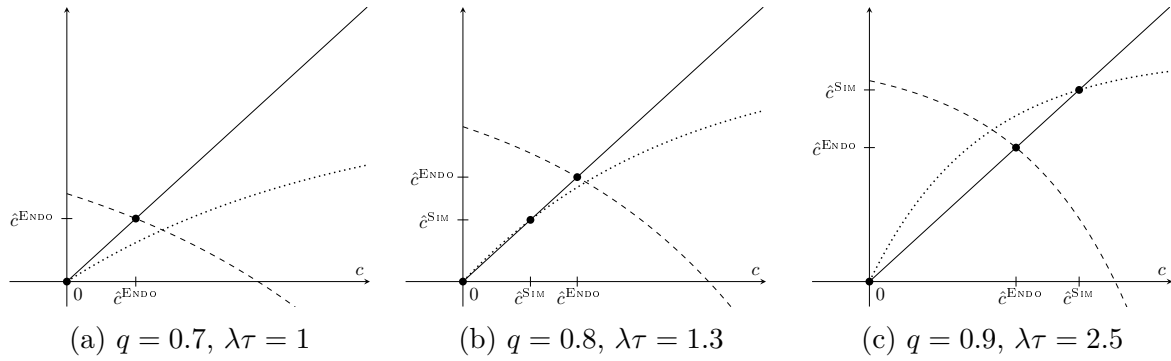
(a) $q = 0.7$, $\lambda\tau = 1$ (b) $q = 0.8$, $\lambda\tau = 1.3$ (c) $q = 0.9$, $\lambda\tau = 2.5$

Figure 3.1: Effect of parameters on equilibrium cutoffs.

The plots represent the following curves: dashed, Eq. (IC$^{\text{Sim}}$); dotted: Eq. (IC$_{\hat{c}}$); solid, 45-degree line.

simultaneous benchmark for difficult crowdfunding campaigns, characterized by a low average bidding $\lambda q \tau$ and a thick upper tail of $F(\cdot)$. In particular, for the binary case, low $\lambda q \tau$ implies $\tau \leq \hat{\tau}(\lambda, c_H)$. Moreover, a thick upper tail of F is equivalent to a large value of $1 - z$, hence the thick tail requirement translates to $z > 1/2$. In the case of a continuum of costs, $\lambda q \tau$ sufficiently low implies $\hat{c}^{\text{Endo}} \geq \hat{c}^{\text{Sim}}$. Also, increasing the thickness of upper tail in a given $F(c_t)$ replacing it with a first-order-statistically-dominant $F'(c_t)$ with $F'(c_t) \leq F(c_t)$ for all c_t makes the sufficient condition (3.35) easier. We have just proved the conclusive proposition of the design section.

Proposition 3.7. *With heterogeneous inspection costs and $g_0 = 2$, $W^{\text{Endo}} \geq W^{\text{Sim}}$ for sufficiently low $\lambda q \tau$ and a sufficiently thick upper tail of $F(\cdot)$.*

3.6   Bid profiles and U-shape

As in Chapter 2, we are also interested in characterizing the temporal profile of bids in Endo and comparing it with Exo.

Since in practice most crowdfunding campaigns are advertised before being officially initiated on the crowdfunding portal, when studying bid profile we augment the bidder population in Endo with a group of pre-arrivals that joins the campaign before it starts and moves at $t = 0$ as soon as the campaign opens. We denote this extension by Endo-Pre. For concreteness, we suppose that the entrepreneur promotes the campaign before its launch on the crowdfunding platform during a period of time τ_0 with constant advertising intensity and attracts bidders' curiosity at the same intensity λ at which bidders discover the campaign once it starts. In that case, the number of pre-arrivals is Poisson distributed with parameter $\Lambda_0 \equiv \lambda \tau_0$. For simplicity, we assume that pre-arrivals do not use the delay option as they find it too costly to evaluate the campaign project at more than one decision episode, but they do delay *implementation* of any decision until

the campaign's opening since that is the first viable moment to do so. In this way, these bidders choose simultaneously $\mathbf{C}$ or $\mathbf{A}$ at time $t = 0$ so that their addition only makes the initial gap stochastic but does not add a positive initial number of delayers.[7] The main analysis extends easily by replacing the initial fixed gap g_0 with its stochastic counterpart $\tilde{g}_0$, which is observed at $t = 0_+$.

The bid profile A_t traces the temporal evolution of the average bids the whole campaign duration. The overall profile can be decomposed into an interim part, for $t \in (0, \tau)$, plus initial and conclusive spikes $\mathbb{A}_0 \triangleq A_{0_+} - A_0$; $\mathbb{A}_\tau \triangleq A_{\tau_+} - A_\tau$ caused by batched arrivals and returns of pre-aware bidders and delayers.

In the interim of the campaign, our equilibria induce a bid rate $\beta_{(t,g)}$ at each funding state (t, g). So A_t is simply the average of those bid rates over possible funding states. In formula,

$$A_t = \sum_{g=-\infty}^{g_0} Q_{(t,g)} \beta_{(t,g)} \tag{3.36}$$

where $Q_{(t,g)}$ is the state-transition probability that the campaign reaches state (t, g) from the initial state $(0, g_0)$, which obeys the differential equation

$$\dot{Q}_{(t,g)} = Q_{(t,g+1)} \beta_{(t,g+1)} - Q_{(t,g)} \beta_{(t,g)}, \qquad t \le \tau \tag{ODE-Q}$$

together with the terminal condition $Q_{(0,g_0)} = 1$. The behavior of the bid profile in the interim part can be studied by analyzing the slope $\dot{A}_t \triangleq \mathrm{d}A/\mathrm{d}t$. It is immediate to see that the size of the initial peak is $\mathbb{A}_0 = \lambda q \tau_0$. Nevertheless, the stochastic produced on g_{0_+} by the exact number of bids pledged pre-arrivals, (not just their average), has non-trivial effects on bid dynamics. The final peak instead has size $\mathbb{A}_\tau = \lambda q \mathbb{E}_{(0,g_0)} (\mu(\mathbf{D}))$ where the expectation is over possible paths of the gap and is taken at the starting point of the campaign.

In Chapter 2 we show that the bid profile in EXO for the homogenous case is decreasing because the campaign either (i) starts frozen, (ii) starts and remains active with a constant bid rate, (iii) starts active but eventually becomes frozen and the bid rate drops to zero. Interestingly, the stationary equilibrium in ENDO leads to the opposite result and can turn the profile from decreasing to increasing. We now show this formally for a $g_0 = 2$ campaign.[8]

In the stationary equilibrium of ENDO-PRE with $g_0 = 2$, the bid rate is $\beta_{(t,2)} = \lambda q \rho$, with ρ in Eq. (3.19) and $\beta_{(t,g)} = \lambda q$ for $g \le 1$. The campaign is in a state $(t, 2)$ if either there is no bid from pre-arrivals, with probability $1 - \sigma(\tau_0)$ and the stopping time of

---

[7]As for SIM, Pareto optimality rules-out a coordination failure from their simultaneous move. There is in any case no such coordination failure risk if their bid-conditional success probability estimated based only on successors is sufficiently high.

[8]If we tie-break with $\mathbf{D}$ rather than $\mathbf{C}$ when $g_t \le 1$ we get a decreasing profile with a final spike. However, it is more reasonable to tie-break with $\mathbf{C}$ since bidders have nothing to learn by waiting and are neutral to other bidders' payoffs.

the first bid pledged after the campaign is started $T_1 \equiv T$ satisfies $T > t$. Since T is exponentially distributed with parameter $\beta_{(t,2)}$ and using the product property of $\sigma(\cdot)$ we have

$$Q_{(t,2)} = \mathbb{P}(g_{0_+} = 2)\mathbb{P}(T > t) = 1 - \sigma(\rho t + \tau_0)$$

This suffices to determine the average bid rate within the interim of the campaign. Precisely,

$$A_t = \lambda q \left[\rho + \sigma(\rho t + \tau_0)(1 - \rho)\right] \tag{3.37}$$

Since $\rho(\cdot)$ is increasing, the profile is positively sloped.[9]

The final spike $\mathbb{A}_\tau$ is the expected bidding of delayers that arrive prior to the realization of T. If there is at least one bid from pre-arrivals, with probability $\sigma(\tau_0)$, the final spike is null. It is also null if given no bid from pre arrivals, $T > \tau$ so that no delayer bids at τ and the campaign fails. If instead $T \leq \tau$, the final spike is positive and given by $\lambda q(1 - \rho)T$, so that $\mathbb{A}_\tau = (1 - \sigma(\tau_0))\lambda q(1 - \rho)\mathbb{E}\left(T \, \mathbb{1}_{\{T \leq \tau\}}\right)$. Computing the last expectation we obtain the expression for $\mathbb{A}_\tau$.

Lemma 3.4. *The expected size of the final bid spike $\mathbb{A}_\tau$ is given by*

$$\mathbb{A}_\tau = (1 - \sigma(\tau_0))S\frac{1 - \rho}{\rho} \tag{3.38}$$

Proof *in Appendix 3.A.*

Expression (3.38) shows τ_0 reduces the final spike, but other parameters have a non-monotone effect. For example, increasing τ reduces the delay probability but also makes it more likely that one bid from **C** bidders motivates **D** bidders to pledge. The effect is positive when τ is raised from a low value but gets negative when it is already large. The same reasoning is valid for q and λ, while increasing c produces a first negative and then positive effect.

With the addition of pre-arrivals, ENDO generates a bimodal profile as a result of the aggregation of two very different patterns. One is decreasing and occurs when there is at least one bid from pre-arrivals, with probability $\sigma(\tau_0)$ that convinces all their successors never to delay. Conditional on this event, we have an initial spike of expected size

$$\mathbb{A}_0 = \lambda q \tau_0 \tag{3.39}$$

a constant bid rate of λq during the campaign, and no final spike.

If instead there is no bid from pre-arrivals, with probability $1 - \sigma(\tau_0)$, the bid profile shows a smooth increase in the bid rate in the interim phase of the campaign and culminates in a final peak.

---

[9]The average profile is concave, but this is a unique feature of $g_0 = 2$

Proposition 3.8. *In* ENDO-PRE*, under homogeneous inspection costs with* $g_0 = 2$*, the bid profile exhibits an initial and final spike* $\mathbb{A}_0$ *and* $\mathbb{A}_\tau$*, given by expressions* (3.39) *and* (3.38)*, together with an increasing bid rate* A_t*, in equation* (3.37)*, during the campaign's interior phase.*

We just proved that the bid profile is binomial in the stationary equilibrium. Nevertheless, due to the equilibrium multiplicity, dynamics are to some degree arbitrary in other non-stationary equilibria. Among these, the final spike is maximal in the p.s.e. with $\mathcal{T}^{\mathbf{D}} = [0, \hat{\tau}(\lambda, c)]$ so that all delays occur initially. In this case, $\mathbb{A}_\tau = (1 - \sigma(\tau_0))\sigma(\tau - \hat{\tau}(\lambda, c))\lambda q\hat{\tau}(\lambda, c)$. Here the spike increases in τ since the probability that **C** bidders convince delayers to bid rises. Conversely, its minimum value of $\mathbb{A}_\tau = 0$ is achieved with all **D** occurring at $[\tau - \hat{\tau}(\lambda, c), \tau]$. The value achieved in other equilibria will be in between those two. The exercise we performed to characterise the bid dynamics under the stationary m.s.e. is important since its properties hold in any stationary equilibria with $g_0 = 2$. Moreover, as we saw in Section 3.4.2, under a continuum of inspection costs, such equilibrium is the unique prediction of the crowdfunding game.

Next, it is worth asking under what conditions the profile is bimodal. To do so, we will first show that in $g_0 = 2$ setting, bimodality is preserved in any non-trivial equilibrium.

That $g_0 = 2$ and heterogeneity always generate a bimodal profile in ENDO-PRE is very intuitive. Pre-arrivals clearly generate the initial peak. Since $F(\hat{c}) \in (0, 1)$, we also have $\mathbb{A}_\tau > 0$. This is clear from the following facts. First, a spike requires some delay and hence no immediate bid when the campaign starts. Ignoring pre-arrivals, this requires no immediate bid during the first infinitesimal interval $(0, \epsilon)$ in which the campaign is open, which has a probability $1 - O(\epsilon)$ by the properties of the Poisson process[10] . It follows that the expected spike is bounded below by the product of $(1 - O(\epsilon))\epsilon F(\hat{c})$, which is the expected number of bids pledged at τ by delayers arriving in the first infinitesimal interval of crowdfunding, *times* the probability of spike activation that is bounded below by $S_{(\epsilon, g_0)} = S_0 - O(\epsilon)$, again by Poisson features. This implies a lower bound $\mathbb{A}_\tau \geq \epsilon F(\hat{c})S_0 + O(\epsilon^2)$, and hence the final spike is positive for small ϵ as S_0 is strictly positive given a non-trivial campaign. Multiplying this lower bound by $1 - \sigma(\Lambda_0)$ adapts to the fact that pre-arrivals may bid and preclude delays. This lower bound is far from tight but it is sufficient to prove the presence of a final spike.

For $g_0 > 2$, we conjecture that heterogenous inspection costs always generate $\mathbb{A}_\tau > 0$ in ENDO and ENDO-PRE. The final spike is only null if bidders never play **D** except once the project freezes, namely at $t : S_{(t,g)} = 0$. However, this scenario with no delays until frozen can only happen for the EXO-equivalent equilibrium with *homogenous* inspection costs because any cost heterogeneity implies that some types play **C** at the same dates at which others would play **D**.

---

[10]$O(\epsilon)$ uses the big-O notation (see Appendix 2.C.1)

3.7   Conclusion

The study of endogenous timing of action in crowdfunding is intriguing but challenging. We provide a preliminary analysis by fully characterizing a threshold two campaign. The stationarity of incentives allows to essentially describe all equilibria by fixing a total number of potential delayers, but this feature is unique to the threshold two setting, making it a very special case. Despite being partial, its analysis already reveals important design features that campaigns with a larger bidder threshold will satisfy. In general, we expect that simultaneity is most useful when costs are homogeneous while endogenous sequencing helps the most when costs are heterogeneous and the upper tail of the cost distribution is thick.

The next logical step towards understanding further endogenous action sequencing in crowdfunding is solving a threshold three campaign. In this setting incentives will not always be stationary any longer. We will also face the additional complication that bidders observing a gap of two would not know precisely the stopping time of the first bid and hence will have to infer it. If costs are homogeneous, bidders will consider that stopping time is uniformly distributed, but this will not be true anymore for bigger initial gaps or heterogenous costs. As a result, the analysis may become analytically intractable, making numerical solutions the only viable way forward. It is also interesting to verify whether certain equilibria such as that with all L-types being decisive and all H-types as delayers will continue to exist in the binary case. In short, the structure of equilibria will undoubtedly change but we expect our main insights to remain relevant.

To respond to the exaggerated control that we give to the platforms in studying design, we also plan to consider a model in which bidders are exogenously divided into exo- and endo-types. EXO bidders are those who only contemplate making all their decisions within the same thinking episode in which they recognize the campaign (e.g. by watching a promotional video or reading the crowdfunding product's description). ENDO-bidders have the same action set as in ENDO.

Finally, if we relax the assumption that bidders find it too costly to follow the project's evolution over time or keep returning to it at regular intervals, we obtain a model where bidders can continuously follow the campaign. In this case, bidders delay by keeping track of the temporal evolution of the campaign until they either inspect or quit. The final spike produced by this model would be smoothed over the whole campaign duration depending on the expected timing at which delaying bidders choose to inspect so that it would predict a different bid profile than ENDO. This setting is more difficult to solve.

Future work on that challenge promises to bear new insights.

Appendices of Chapter 3

3.A Proofs

Proof of Lemma 3.1.   The bidder surplus V can be tackled by calculating separately and then combining the surplus of bidders arriving during the time intervals $[0, \tau - \hat{\tau}(c, \lambda)]$ and $[\tau - \hat{\tau}(c, \lambda), \tau]$. In order to proceed, we use the property of Poisson arrivals such that, conditional on a bidder arriving during the campaign, his exact arrival time is uniformly distributed over $[0, \tau]$. So we apply the decomposition

$$V^{\text{Exo}} = \frac{\tau - \hat{\tau}(\lambda, c)}{\tau} V_{t \in [0, \tau - \hat{\tau}(\lambda, c)]} + \frac{\hat{\tau}(\lambda, c)}{\tau} V_{t \in [\tau - \hat{\tau}(c, \lambda), \tau]}$$

In the first interval, all bidders always play **C**. They obtain a payoff $(q - c)$ if $g_t \leq 1$ and $q\sigma(\tau - t) - c$ for $g_t = 2$. Since the probability of no bid by t is the $1 - \sigma(t)$ and $(1 - \sigma(t))(1 - \sigma(\tau - t)) = 1 - \sigma(\tau)$, the average bidder surplus in the first time interval is

$$V_{t \in [0, \tau - \hat{\tau}(\lambda, c)]} = q\sigma(\tau) - c$$

For the remaining duration $t \in [\tau - \hat{\tau}(c, \lambda), \tau]$, bidders inspect only for $g_t \leq 1$, with probability $\sigma(\tau - \hat{\tau}(\lambda, c))$, and get $U^{\mathbf{C}}_{(t,1)} = q - c$. So, using the definition of $\sigma(\cdot)$ in Eq. (3.8) we can again see that

$$V_{t \in [\tau - \hat{\tau}(c, \lambda), \tau]} = (q - c)\sigma(\tau - \hat{\tau}(\lambda, c)) = (q - c)\left(1 - e^{-\lambda q \tau} \frac{q}{q - c}\right)$$

$$= (q - c) - qe^{-\lambda q \tau} = q\sigma(\tau) - c$$

The overall expected surplus in (3.13) coincides with the identical value achieved in each time region. ∎

Proof of Lemma 3.2.   We proceed by showing that the payoff difference $U^{\mathbf{C}}_{(t,2)} - U^{\mathbf{D}}_{(t,2)}$ is proportional to $1 - \phi_{(t,2)}$. In other words, the incentive constraint is evaluated in a scenario in which $g_\tau = 2$ so a free-riding intent by t would fail since only in that case the payoffs of a decisive bidder and a delayer eventually differ. After factoring-out $1 - \phi_{(t,2)}$, the incentive constraint becomes stationary and depends only on a bidder's pivotality on the campaign in state $(\tau, 2)$, hence only on the total mass of potential delayers δ.

Concretely, subtracting the expected utility values in Eqs. (3.17) and (3.18) we obtain

$$
\begin{aligned}
U^{\mathbf{C}}_{(t,2)} - U^{\mathbf{D}}_{(t,2)} &= q\left[\sigma\left(\kappa_t + \delta^{\mathrm{pre}}_t\right) - \sigma\left(\kappa_t - \delta^{\mathrm{post}}_t\right)\right] - c\left(1 - \sigma\left(\kappa_t - \delta^{\mathrm{post}}_t\right)\right) \\
&= q\left[\left(1 - \sigma\left(\kappa_t - \delta^{\mathrm{post}}_t\right)\right) - \left(1 - \sigma\left(\kappa_t + \delta^{\mathrm{pre}}_t\right)\right)\right] - c\left(1 - \sigma\left(\kappa_t - \delta^{\mathrm{post}}_t\right)\right) \\
&= \left(1 - \sigma\left(\kappa_t - \delta^{\mathrm{post}}_t\right)\right)\left\{q\left[1 - \left(1 - \sigma\left(\kappa_t + \delta^{\mathrm{pre}}_t\right)\right)\left(1 - \sigma\left(\kappa_t - \delta^{\mathrm{post}}_t\right)\right)^{-1}\right] - c\right\}
\end{aligned}
$$

$$(3.A.1)$$

Using the properties of exponentials and given that $1 - \sigma(x) = e^{-x}$, we have that

$$
\left(1 - \sigma\left(\kappa_t + \delta^{\mathrm{pre}}_t\right)\right)\left(1 - \sigma\left(\kappa_t - \delta^{\mathrm{post}}_t\right)\right)^{-1} = 1 - \sigma(\delta), \quad \text{with} \quad \delta = \delta^{\mathrm{pre}}_t + \delta^{\mathrm{post}}_t
$$

So, after substituting $\left(1 - \sigma\left(\kappa_t - \delta^{\mathrm{post}}_t\right)\right)$ with $1 - \phi_{(t,2)}$, equation (3.A.1) simplifies to

$$
\phi_{(t,2)}\left[q\sigma(\delta) - c\right]
$$

The inequality $U^{\mathbf{C}}_{(t,2)} - U^{\mathbf{D}}_{(t,2)} \geq 0$ is determined by the above expression in squared brackets which coincides with the incentive-compatibility constraint displayed in $(\mathrm{IC}^{\mathbf{C,D}}_c)$. ∎

Proof of Proposition 3.4.   Stationary of equilibrium strategies follows from the fact that $(\mathrm{IC}_{\hat{c}})$ has to hold at every instant.

To prove that $\hat{c}$ is interior and unique, first, we have that $\hat{c} = 0$ violates $(\mathrm{IC}_{\hat{c}})$ since $\sigma(\tau(1 - F(0))) = \sigma(\tau) > 0$, so inactivity with $\hat{c} = 0$ is not an equilibrium. We also cannot an equilibrium with $\hat{c} = q$ as it again violates $(\mathrm{IC}_{\hat{c}})$ as $\sigma(\tau(1 - F(q))) = 0$ and so $\sigma(\tau(1 - F(q))) - 1 < 0$. We can conclude that $\hat{c} \in (0, q)$.

Furthermore, the left-hand side of Eq. $(\mathrm{IC}_{\hat{c}})$ is continuous and monotone decreasing in $\hat{c}$, as we can confirm by taking its partial derivative with respect to $\hat{c}$ and see that it is always negative:

$$
-\left[\lambda q\tau f(\hat{c})(1 - \sigma(\tau(1 - F(\hat{c})))) + 1/q\right] < 0 \quad \text{for all } \hat{c}
$$

If follows that there is exactly one value $\hat{c} \in (0, q)$ that satisfies Eq. $(\mathrm{IC}_{\hat{c}})$. ∎

Proof of Proposition 3.5.   The first- and second- order conditions to problem (3.31) are as follows

$$
\sigma(\tau_1)(q - c)(1 + \lambda q(\tau - \tau_1)) = q\sigma(\tau) - c + (q - c)\lambda q(\tau - \tau_1) \tag{FOC}
$$

$$
-(q - c)(1 - \sigma(\tau_1))(2 + \lambda q(\tau - \tau_1)) \leq 0 \tag{SOC}
$$

At this point, we can see that, (FOC) is violated at $\tau_1 = \tau$, because $\sigma(\tau)(q-c) - (q\sigma(\tau) - c) > 0$. This proves $V^{\text{SIM}} < V^*$ which is achieved at a value $\hat{\tau}^* > 0$. For $\tau_1 = \tau - \hat{\tau}(\lambda, c)$ condition (FOC) is again violated because using the surplus-equivalence result in Proposition 3.2 combined with V in Eq. (3.10),

$$\sigma(\tau_1)(q-c)(1 + \lambda q(\tau - \tau_1)) = q\sigma(\tau) - c + \sigma(\tau_1)(q-c)\lambda q(\tau - \tau_1)$$
$$< q\sigma(\tau) - c + (q-c)\lambda q(\tau - \tau_1)$$

So of course $V^{\text{ENDO}} = V^{\text{SIM}} < V^*$ and $\hat{\tau}^* < \hat{\tau}(\lambda, c)$. This proves that any value $\hat{\tau} \in (0, \hat{\tau}(\lambda, c))$ raises bidder surplus relative to SIM and ENDO. ∎

Proof of Proposition 3.6. ENDO prescribes inactivity $\tau < \hat{\tau}(\lambda, c_L)$. Since $\hat{\tau}(\lambda, c_L) < \hat{\tau}(\lambda z, c_L)$, where the latter is the condition for inactivity in SIM, we have $V^{\text{ENDO}} \geq V^{\text{SIM}} = 0$ for this range of values.

For $\tau \in [\hat{\tau}(\lambda z, c_L), \hat{\tau}(\lambda, c_H))$, ENDO makes both H-type- and L-types better-off since $V_H > 0$ and $zV_L > V^{\text{SIM}}$ for V_L in Eq. (3.26) and V^{SIM} Eq. (3.33) as the former has $\sigma(\tau)$ while the latter only $\sigma(z\tau)$. Therefore $V^{\text{ENDO}} \geq V^{\text{SIM}}$.

For $\tau \geq \hat{\tau}(\lambda, c_H)$, L-types are indifferent between ENDO and SIM but H-types are either indifferent or better-off because their surplus in ENDO is $V_H = (q - c_H)\sigma(\tau - \delta) \geq q\sigma(\tau) - c_H$ as $\delta \leq \hat{\tau}(\lambda, c_H)$.

The proof of the comparison of S_0 values is provided in the paragraph preceding the proposition statement and the overall comparison of total welfare is immediate combining its surplus and success rate components ∎

Proof of Lemma 3.4.

$$\mathbb{E}\left(T \, \mathbb{1}_{\{T \leq \tau\}}\right) = \int_0^\tau e^{-\lambda q \rho T} \lambda q \rho T \, dT \qquad = \left[e^{-\lambda q \rho T}\left(\frac{\lambda q \rho T + 1}{\lambda q \rho}\right)\right]_{T=\tau}^{T=0}$$

$$= \frac{1}{\lambda q \rho} - e^{-\lambda q \rho \tau}\left(\tau + \frac{1}{\lambda q \rho}\right) \qquad = \frac{1 - e^{-\lambda q \rho \tau} - \lambda q \rho \tau e^{-\lambda q \rho \tau}}{\lambda q \rho}$$

so that $(1 - \sigma(\tau_0))\lambda q(1 - \rho)\mathbb{E}\left(T \, \mathbb{1}_{\{T \leq \tau\}}\right) = \mathbb{A}_\tau = (1 - \sigma(\tau_0))S\dfrac{1 - \rho}{\rho}$

∎

Page intentionally left blank.

Bibliography

Abadi, J. and Brunnermeier, M. (2018). Blockchain economics. Technical report, National Bureau of Economic Research.

Admati, A. R. and Perry, M. (1991). Joint projects without commitment. *Review of Economic Studies*, 58(2):259–276.

Agrawal, A., Catalini, C., and Goldfarb, A. (2014). Some simple economics of crowdfunding. *Innovation Policy and the Economy*, 14(1):63–97.

Agrawal, A. K., Catalini, C., and Goldfarb, A. (2011). The geography of crowdfunding. Technical report, National bureau of economic research.

Alaei, S., Malekian, A., and Mostagir, M. (2016). A dynamic model of crowdfunding.

Athey, S., Parashkevov, I., Sarukkai, V., and Xia, J. (2016). Bitcoin pricing, adoption, and usage: Theory and evidence.

Babich, V., Tsoukalas, G., and Marinesi, S. (2017). Does crowdfunding benefit entrepreneurs and venture capital investors?

Belleflamme, P., Lambert, T., and Schwienbacher, A. (2014). Crowdfunding: Tapping the right crowd. *Journal of business venturing*, 29(5):585–609.

Biais, B., Bisiere, C., Bouvard, M., and Casamatta, C. (2019). The blockchain folk theorem. *The Review of Financial Studies*, 32(5):1662–1715.

Budish, E. (2018). The economic limits of bitcoin and the blockchain. Technical report, National Bureau of Economic Research.

Carlsten, M., Kalodner, H., Weinberg, S. M., and Narayanan, A. (2016). On the instability of bitcoin without the block reward. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 154–167.

Cason, T. N. and Zubrickas, R. (2018). Crowdfunding for public goods with refund bonuses: An empirical and theoretical investigation.

Chakraborty, S. and Swinney, R. (2019). Designing rewards-based crowdfunding campaigns for strategic (but distracted) contributors. *Available at SSRN 3240094*.

Chang, J.-W. (2016). The economics of crowdfunding.

Chemla, G. and Tinn, K. (2018). Learning through crowdfunding. *Available at SSRN 2796435*.

Chiu, J. and Koeppl, T. (2019). The Economics of Cryptocurrencies—Bitcoin and Beyond. Staff Working Papers 19-40, Bank of Canada.

Choi, M. and Rocheteau, G. (2020a). Money mining and price dynamics. *Available at SSRN 3336367*.

Choi, M. and Rocheteau, G. (2020b). More on money mining and price dynamics: Competing and divisible currencies. *Available at SSRN*.

Choi, M. and Rocheteau, G. (2020c). New monetarism in continuous time: Methods and applications. *Available at SSRN 3435889*.

Colombo, M. G., Franzoni, C., and Rossi-Lamastra, C. (2015). Internal social capital and the attraction of early contributions in crowdfunding. *Entrepreneurship theory and practice*, 39(1):75–100.

Cong, L. W., He, Z., and Li, J. (2019). Decentralized mining in centralized pools. Technical report, National Bureau of Economic Research.

Cordova, A., Dolci, J., and Gianfrate, G. (2015). The determinants of crowdfunding success: evidence from technology projects. *Procedia-Social and Behavioral Sciences*, 181:115–124.

Crosetto, P. and Regner, T. (2018). It's never too late: funding dynamics and self pledges in reward-based crowdfunding. *Research Policy*, 47(8):1463–1477.

Deb, J., Öry, A., and Williams, K. (2021). Aiming for the goal : Contribution dynamics of crowdfunding. *Cowles Foundation Discussion Papers. 2597*.

Decker, C. and Wattenhofer, R. (2013). Information propagation in the bitcoin network. In *IEEE P2P 2013 Proceedings*, pages 1–10. IEEE.

Easley, D., O'Hara, M., and Basu, S. (2019). From mining to markets: The evolution of bitcoin transaction fees. *Journal of Financial Economics*.

Ellman, M. and Fabi, M. (2021a). Crowdfunding with endogenously timed moves.

Ellman, M. and Fabi, M. (2021b). A theory of crowdfunding dynamics. *Barcelona GSE working paper*.

Ellman, M. and Hurkens, S. (2019a). Fraud tolerance in optimal crowdfunding. *Economics Letters*, 181(C):11–16.

Ellman, M. and Hurkens, S. (2019b). Optimal crowdfunding design. *Journal of Economic Theory*, 184:104939.

Ennis, H. M. (2009). Avoiding the inflation tax. *International Economic Review*, 50(2):607–625.

Etter, V., Grossglauser, M., and Thiran, P. (2013). Launch hard or go home! predicting the success of kickstarter campaigns. In *Proceedings of the first ACM conference on Online social networks*, pages 177–182.

Eyal, I. and Sirer, E. G. (2014). Majority is not enough: Bitcoin mining is vulnerable. In *International conference on financial cryptography and data security*, pages 436–454. Springer.

Fernández-Villaverde, J. and Sanches, D. (2019). Can currency competition work? *Journal of Monetary Economics*, 106:1–15.

Goldfarb, A. and Tucker, C. (2019). Digital economics. *Journal of Economic Literature*, 57(1):3–43.

Grunspan, C. and Pérez-Marco, R. (2018). Double spend races. *International Journal of Theoretical and Applied Finance*, 21(08):1850053.

Houy, N. (2016). The bitcoin mining game. *Ledger*, 1:53–68.

Hu, M., Li, X., and Shi, M. (2015). Product and pricing decisions in crowdfunding. *Marketing Science*, 34(3):331–345.

Huberman, G., Leshno, J., and Moallemi, C. C. (2019). An economic analysis of the bitcoin payment system. *Columbia Business School Research Paper*, (17-92).

Kroll, J. A., Davey, I. C., and Felten, E. W. (2013). The economics of bitcoin mining, or bitcoin in the presence of adversaries. In *Proceedings of WEIS*, volume 2013, page 11.

Kuppuswamy, V. and Bayus, B. L. (2015). Crowdfunding creative ideas: The dynamics of project backers in kickstarter.

Kuppuswamy, V. and Bayus, B. L. (2017). Crowdfunding creative ideas: The dynamics of project backers in kickstarter. *A shorter version of this paper is in" The Economics of*

Crowdfunding: Startups, Portals, and Investor Behavior"-L. Hornuf and D. Cumming (eds.).

Kuppuswamy, V. and Bayus, B. L. (2018). 17. a review of crowdfunding research and findings. *Handbook of Research on New Product Development, Golder, PN, and Mitra, D.(Eds.), Cheltenham, UK: Edward Elgar Publishing*, pages 361–373.

Lagos, R., Rocheteau, G., and Wright, R. (2014). The art of monetary theory: A new monetarist perspective. *forthcoming, Journal of Economic Literature.*

Lagos, R. and Wright, R. (2005). A unified framework for monetary theory and policy analysis. *Journal of political Economy*, 113(3):463–484.

Leshno, J. D. and Strack, P. (2019). Bitcoin: An axiomatic approach and an impossibility theorem. *American Economic Review: Insights.*

Liu, S. (2020). A theory of collective investment with application to venture funding.

Ma, C.-T. A. and Manove, M. (1993). Bargaining with Deadlines and Imperfect Player Control. *Econometrica*, 61(6):1313–1339.

Mollick, E. (2014). The dynamics of crowdfunding: An exploratory study. *Journal of business venturing*, 29(1):1–16.

Myerson, R. B. (1978). Refinements of the nash equilibrium concept. *International Journal of Game Theory*, 7:73–80.

Myerson, R. B. (1998). Population uncertainty and poisson games. *International Journal of Game Theory*, 27(3):375–392.

Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. *Cryptography Mailing list at https://metzdowd.com.*

Neudecker, T. and Hartenstein, H. (2019). Short paper: An empirical analysis of blockchain forks in bitcoin. In *International Conference on Financial Cryptography and Data Security*, pages 84–92. Springer.

Øksendal, B. K. and Sulem, A. (2007). *Applied stochastic control of jump diffusions*, volume 498. Springer.

Pinzón, C. and Rocha, C. (2016). Double-spend attack models with time advantange for bitcoin. *Electronic Notes in Theoretical Computer Science*, 329:79–103.

Prat, J. and Walter, B. (2018). An equilibrium model of the market for bitcoin mining.

Rao, H., Xu, A., Yang, X., and Fu, W.-T. (2014). Emerging dynamics in crowdfunding campaigns. In *International Conference on Social Computing, Behavioral-Cultural Modeling, and Prediction*, pages 333–340. Springer.

Rizun, P. R. (2015). A transaction fee market exists without a block size limit. *Block Size Limit Debate Working Paper*.

Rosenfeld, M. (2014). Analysis of hashrate-based double spending. *arXiv preprint arXiv:1402.2009*.

Rosu, I. and Saleh, F. (2019). Evolution of shares in a proof-of-stake cryptocurrency. *Available at SSRN 3377136*.

Saad, M., Spaulding, J., Njilla, L., Kamhoua, C., Shetty, S., Nyang, D., and Mohaisen, A. (2019). Exploring the attack surface of blockchain: A systematic overview. *arXiv preprint arXiv:1904.03487*.

Sahm, M. (2016). Advance-purchase financing of projects with few buyers.

Saleh, F. (2020). Blockchain without waste: Proof-of-stake. *Available at SSRN 3183935*.

Schilling, L. and Uhlig, H. (2019). Some simple bitcoin economics. *Journal of Monetary Economics*, 106:16–26.

Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, 4:25–55.

Simon, L. and Stinchcombe, M. (1995). Equilibrium refinement for infinite normal-form games. *Econometrica*, 63(6):1421–43.

Solomon, J., Ma, W., and Wash, R. (2015). Don't wait!: How timing affects coordination of crowdfunding donations. In *Proceedings of the 18th acm conference on computer supported cooperative work & social computing*, pages 547–556. ACM.

Strausz, R. (2017). A theory of crowdfunding: A mechanism design approach with demand uncertainty and moral hazard. *American Economic Review*, 107(6):1430–76.

Van Damme, E. (1991). *Stability and perfection of Nash equilibria*. Springer-Verlag, Berlin. Cota antiga E.130.073.

Vismara, S. (2016). Information cascades among investors in equity crowdfunding. *Entrepreneurship Theory and Practice*.

Zhang, J. (1997). Strategic delay and the onset of investment cascades. *The RAND Journal of Economics*, pages 188–205.

Zhang, J. and Liu, P. (2012). Rational herding in microloan markets. *Management science*, 58(5):892–912.

Zhang, J., Savin, S., and Veeraraghavan, S. (2017). Revenue management in crowdfunding.