Non-Clairvoyant Dynamic Mechanism Design*†

Vahab Mirrokni Renato Paes Leme Pingzhong Tang Song Zuo

October 13, 2018

Abstract

Despite their better revenue and welfare guarantees for repeated auctions, dynamic mechanisms have not been widely adopted in practice. This is partly due to the complexity of their implementation as well as their unrealistic use of forecasting for future periods. We address these shortcomings and present a new family of dynamic mechanisms that are simple and require no distributional knowledge of future periods.

This paper introduces the concept of non-clairvoyance in dynamic mechanism design, which is a measure-theoretic restriction on the information that the seller is allowed to use. A dynamic mechanism is non-clairvoyant if the allocation and pricing rule at each period does not depend on the type distributions in the future periods.

We develop a framework (bank account mechanisms) for characterizing, designing, and proving lower bounds for dynamic mechanisms (clairvoyant or non-clairvoyant). This framework is used to characterize the revenue extraction power of the non-clairvoyant mechanisms with respect to the mechanisms that are allowed unrestricted use of distributional knowledge.

*Thanks are due to Dirk Bergemann for his detailed comments on an early draft of this manuscript. We also thank Ilan Lobel, Santiago Balseiro and Balu Sivan, the participants of the Erice Workshop of Stochastic Methods in Game Theory and the Google Market Algorithms Workshop for their comments and suggestions.

†Contact information: Vahab Mirrokni, Google Research NYC, `mirrokni@google.com`; Renato Paes Leme, Google Research NYC, `renatoppl@google.com`; Pingzhong Tang, Institute for Interdisciplinary Information Sciences, Tsinghua University, `kenshinping@gmail.com`; Song Zuo, Institute for Interdisciplinary Information Sciences, Tsinghua University, `songzuo.z@gmail.com`.

1 Introduction

Dynamic mechanism design is a powerful tool for designing auctions that are repeated over time. Optimizing auctions across different time periods instead of optimizing each period individually can lead to improvements, both in terms of revenue and allocation efficiency. So, why isn't the adoption of dynamic mechanisms more widespread? A crucial limitation is that they tend to be too detail-dependent: they require the designer to have reliable forecasts of the distributions of the agent's valuations in all periods. Moreover, the mechanisms often require all the buyers to share those beliefs. A second reason is their descriptive complexity, which leads to non-intuitive allocation and pricing rules. Our goal in this paper is to design mechanisms that address the points raised above.

We will consider the design of sequential auctions. A firm wants to sell T products over T periods of time, one for each period. Each product has features that can be publicly observed by all agents (the seller and the buyers) as well as features that can only be privately observed by each buyer. The valuation of each buyer is a function of the publicly observable features and her privately observable features. In the classic Bayesian formulation of this problem, for each period t, the publicly observable feature vector induces a common-knowledge distribution F_t over the buyer's valuations. Even when the valuation distributions $F_1 \dots F_T$ are independent, the incentive constraints bind across periods. Jackson and Sonnenschein [JS07] observed that linking apparently independent decisions together could lead to improved outcomes. A similar phenomenon was also observed by Manelli and Vincent [MV07]. Gains can be obtained, for example, by offering the buyer a higher price today in exchange for a discount tomorrow. There is a practice in the industry of offering discounts to the buyers that have previously purchased their products. Papadimitriou et al. [PPPR16] quantified the revenue gap between applying the optimal Myerson auction to sell the item in each period and using the optimal mechanism that links decisions across time. They showed that the latter could obtain arbitrarily times more revenue.

The central theme of this paper is to robustify the design of dynamic mechanisms by making them independent of the distributional beliefs about the future. This is achieved by replacing the fully-Bayesian uncertainty in the classic dynamic mechanism design with a mixed uncertainty model, where both the seller and the buyers have Bayesian uncertainty about the present, but Knightian uncertainty about the future periods. Before going into details about our mixed uncertainty model, we describe a concrete economic scenario we wish to analyze and argue that the fully-Bayesian modeling is not robust enough for practical applications.

Motivation and economic application We consider the problem of designing repeated auctions for display internet advertising, which corresponds to the online ads appearing in web-pages such as news sites, blogs, etc. Each time a pageview is issued, the advertising platform sends a real-time bid request to all the relevant advertisers containing the features of the pageview, such as geolocation, device, browser, the page from which the pageview is originating, and importantly, a unique identifier for the user (UserId) issuing the pageview. All this information is public to both the platform and the advertisers have mechanisms to independently verify it.

Advertisers also have private information about that specific UserId. If the advertiser is an online store, it typically stores a database of the UserIds that visited the store website recently. Upon receiving the bid request, the advertiser can check whether that UserId visited the website recently or not and bid accordingly. This mechanism is called *remarketing* and is one of the main sources of information asymmetry between the advertisers and the platform.

Challenges in designing robust mechanisms Properly linking decisions across time requires the precise knowledge of the distributions in each period. A seller trying to deploy a traditional fully-Bayesian mechanism for the problem above would require the knowledge of the buyer distributions in each future period, which would require knowing the public features of the future pageviews in advance.

Forecasting the features of future queries is a hard problem given the non-stochasticity of internet traffic patterns.This is particularly worrisome for fully-Bayesian mechanisms, since the (dynamic) incentive compatibility itself depends on accurate forecasts. Dynamic incentive compatibility means that agents in each period are incentivized to report their types based on their knowledge of the current types and in expectation over their types in the future. So if neither the seller nor the agents can build good forecasts, how can they even verify that a mechanism is dynamic incentive compatible?

Non-Clairvoyance Since the buyer's value distribution F_t only becomes available at time t once the publicly observable features of the product are revealed, a natural class of mechanisms is one where the allocation and pricing cannot depend on the distributional information about the future. We call such mechanisms *non-clairvoyant*. Formally, a non-clairvoyant mechanism consists of an allocation and pricing rule that, for each period t, maps the distributions $F_1, \ldots, F_t$ and the types $\theta_1, \ldots, \theta_t$ sampled from those distributions to an allocation and payments. What does it mean for non-clairvoyant mechanisms to be dynamic incentive compatible? In traditional mechanism design, dynamic incentive compatibility means that it is the optimal strategy for the agent to report her current type truthfully *in expectation* over her types in future periods. We say that a mechanism is dynamic incentive compatible in the non-clairvoyant sense if, for *any* continuation future $F_{t+1}, \ldots, F_T$, it is incentive compatible for the buyer to report her type in period t. This is quite a strong notion, since we do not even require the agents and the designer to agree on the forecast for future periods. Besides dynamic incentive compatibility, we will also require ex-post individual rationality, i.e., the utility of the agent for the mechanism is non-negative for every realization of her types.

To understand the relative power of *non-clairvoyant* and *clairvoyant* mechanisms (i.e., the mechanisms that know all the distributions $F_1, \ldots, F_T$) we must consider two scenarios,

$$\text{scenario } A : F_1, \ldots, F_t, F'_{t+1}, \ldots, F'_T \qquad \text{scenario } B : F_1, \ldots, F_t, F''_{t+1}, \ldots, F''_T.$$

The non-clairvoyant designer needs to design a mechanism that allocates in the same way for the first t periods in both scenarios. The clairvoyant designer can tailor his allocation and payments in the first t periods to his knowledge of whether he is in scenario A or B.

Importance of ex-post participation constraints Besides the mixed model of uncertainty, our second departure from the classic theory of dynamic mechanism design is that we impose ex-post participation constraints instead of the more usual interim constraints. The first reason for such a choice is again based on the online ads application. In such systems, the advertiser is asked to report his/her *maximum willingness to pay* for each ad click or impression and the advertising platform is legally obligated to charge for a bundle of displayed ads at most the maximum value declared by the advertiser for that bundle.

All internet advertising contracts contain some mechanism in which buyers can specify a hard limit on the maximum amount to pay for any given item. This is very important from a business perspective since it allows the advertisers to control their risk.

A second reason for the choice of ex-post constraints is that they are distribution-independent, which adds another layer of robustness to the mechanism. This is opposed to interim constraints, which are taken in expectation over the type distributions.

Techniques The main technique used in the paper both to design non-clairvoyant mechanisms and to upper bound the revenue of any clairvoyant mechanism is a framework which we call *bank account mechanisms*. First, we show that for any clairvoyant mechanism that is dynamic incentive compatible and ex-post individually rational, there is a bank account mechanism with the same properties and at least the same revenue (Lemma 3.2). Bank account mechanisms have several nice properties. First, any mechanism in that format is dynamic incentive compatible by design (Lemma 3.1). Second, its revenue naturally decomposes into the intra-period revenue, which can be bounded by the Myerson revenue for that period, and the inter-period revenue, which we call *bank account spend* (Lemma 4.1). Perhaps more importantly, bank account mechanisms naturally lend themselves to the design of non-clairvoyant mechanisms. This statement is formalized in the following non-clairvoyant reduction: any non-clairvoyant dynamic mechanism can be written as a non-clairvoyant bank account mechanism (Lemma 5.2).

As a result, the bank account mechanisms characterize a class of clairvoyant (or non-clairvoyant) dynamic mechanisms with simple structures, and it is without loss of generality to focus on this class while designing revenue-optimal and/or welfare-efficient dynamic mechanisms.

Main Results Using the techniques we develop, we characterize the optimal non-clairvoyant mechanism for selling one item per period for two periods to multiple buyers (Theorem 6.3). The mechanism is optimal in the sense that it obtains 1/2 of the optimal clairvoyant revenue and 1/2 is the best achievable ratio. We present an impossibility result (Theorem 5.1) showing that no non-clairvoyant mechanism can obtain a better-than-1/2 fraction of the revenue of the optimal clairvoyant mechanism for all sequences of distributions.

The result described above is a special case of a more general construction that holds for any number of periods (Theorem 6.1). We obtain a non-clairvoyant mechanism for selling one item per period for any number of periods, which we call the NonClairvoyantBalance mechanism, that obtains at least 1/5 of the revenue achieveable by the optimal clairvoyant mechanism. Since the revenue of the optimal dynamic mechanism can be arbitrarily times more than the static one, obtaining at least 1/5 of the optimal dynamic revenue often means obtaining much larger revenue than the optimal static auction.

The mechanism sells in each period 1/5 of the item using the Myerson auction for that distribution in that period and 2/5 of the item as a plain second price auction. The remaining 2/5 of the item will provide the dynamic component of the mechanism: for the remaining 2/5, we will use a parameter b^i called *bank balance* computed for each agent as a function of her previous reports and the previous distributions; then we will run a modification of the optimal money burning auction of Hartline and Roughgarden [HR08]. The Myerson auction component will capture the revenue that can be obtained within each individual period. The combination of the second price and the money burning components will be responsible for capturing the gains from inter-period interactions.

2 Repeated Auctions Model

Auction Setup The standard dynamic mechanism design setting with finite time horizon describes an economic setup where a designer repeatedly selects an outcome over T periods based on the reports by strategic agents. For the sake of clarity, the first part of our paper focuses on the single agent case and then extends it to multiple agents in Section 6. In each period $t \in [T]$, the agent has type $\theta_t \in \Theta$, which is drawn from distribution F_t independent across timesteps. Her valuation for outcome $x_t \in \mathcal{O}$ being implemented is given by $v : \Theta \times \mathcal{O} \to \mathbb{R}$.

Our assumption that the agent types are independent across timesteps is inspired by our main application in internet advertising: each time an impression arrives, the advertiser's value is a function of the properties of the impression that are publicly observable (e.g., as geographic and demographic information) plus some privately observable information (e.g., browser cookies).[1] The publicly observable information will determine the distribution from which the buyer's type is sampled, while the private information will determine the realization of the type. Unless buyers are starting a new campaign, they already have an established notion of value for the combination of cookie and demographic information, so the allocation for one impression will not affect the value of others. We also consider implementing dynamic mechanisms with a short span (say a few hours or a day) in which there is a large enough volume of queries that we can reap the benefit of dynamic queries but the time span is short enough for the valuations to remain stable. Concerns about valuations that shift over time arise when we try to apply dynamic mechanisms over large time spans when the market is likely to move. This issue, however, lies beyond the scope of the current paper.

Continuing the description of the model, the following events happen at each period t:

1. the agent learns her type $\theta_t \sim F_t$;
2. the agent reports type $\hat{\theta}_t$ to the designer;
3. the designer implements an outcome $x_t \in \mathcal{O}$ and charges the agent p_t; and
4. the agent accrues utility $u_t = v(\theta_t, x_t) - p_t$.

The final utility of the agent is the sum over her utility in all periods, i.e., $\sum_{t=1}^{T} u_t$.

A mechanism can be described in terms of a pricing and an outcome function, which map the distributional knowledge of the seller $F_{1..T} = (F_1, F_2, \ldots, F_T)$ and the history of reports $\hat{\theta}_{1..t} = (\hat{\theta}_1, \hat{\theta}_2, \ldots, \hat{\theta}_t)$ to an outcome x_t and payment p_t:

- Outcome: $x_t : \Theta^t \times (\Delta\Theta)^T \to \mathcal{O}$,
- Payment: $p_t : \Theta^t \times (\Delta\Theta)^T \to \mathbb{R}$,

[1]Browser cookies are small pieces of data sent from websites and stored in the user's Web browser so that when the user revisits the same website the cookies can be used to identify the user's previous actions. Such data is encrypted and can only be read by the website that placed it. For example, if a user visits an online store, a cookie is placed on his browser. In an auction, only the advertiser corresponding to that store will be able to read that cookie, making it a private signal.

where Θ is the space of types for the agents and $\Delta\Theta$ is a set of distributions over Θ. We use a semicolon to separate the report and distribution parameters: $x_t(\hat{\theta}_{1..t}; F_{1..T})$ and $p_t(\hat{\theta}_{1..t}; F_{1..T})$. We will omit the distributional parameters $F_{1..T}$ when clear from the context and write it simply as $x_t(\hat{\theta}_{1..t})$ and $p_t(\hat{\theta}_{1..t})$.

We define the utility of the buyer with type θ_t in step t given a history of reports $\hat{\theta}_{1..t}$ and seller distribution knowledge $F_{1..T}$ as:

$$u_t(\theta_t; \hat{\theta}_{1..t}; F_{1..T}) = v(\theta_t, x_t(\hat{\theta}_{1..t}; F_{1..T})) - p_t(\hat{\theta}_{1..t}; F_{1..T}),$$

and again we omit $F_{1..T}$ when clear from context.

Incentive Constraints We will adopt the traditional notion of incentive compatibility in dynamic settings, where agents have incentives to report their types truthfully in each period. This can be defined easily by backward induction: in the last period, regardless of the history so far, it should be incentive compatible for an agent to report her true type. This corresponds to the usual notion of incentive compatibility in (static) mechanism design:

$$\theta_T = \arg\max_{\hat{\theta}_T} u_T(\theta_T; \hat{\theta}_{1..T-1}, \hat{\theta}_T; F_{1..T})$$

for all $\hat{\theta}_{1..T-1}, \theta_T$. To simplify notations, from now on we will omit the 'for-all' quantification and assume that all expressions are quantified as 'for-all' in their free variables.

For the next-to-last-period, it should be incentive compatible for the agent to report her true type given that she will report her true type in the following period:

$$\theta_{T-1} = \arg\max_{\hat{\theta}_{T-1}} u_{T-1}(\theta_{T-1}; \hat{\theta}_{1..T-2}, \hat{\theta}_{T-1}; F_{1..T}) + \mathbb{E}_{\theta_T \sim F_T}[u_T(\theta_T; \hat{\theta}_{1..T-2}, \hat{\theta}_{T-1}, \theta_T; F_{1..T})].$$

Proceeding by backward induction for all periods, we require that:

$$\theta_t = \arg\max_{\hat{\theta}_t} u_t(\theta_t; \hat{\theta}_{1..t-1}, \hat{\theta}_t; F_{1..T}) + U_t(\hat{\theta}_{1..t-1}, \hat{\theta}_t; F_{1..T}), \tag{DIC}$$

where the second term is the *continuation utility*, i.e., the expected utility obtained from the subsequent periods of the mechanism:

$$U_t(\hat{\theta}_{1..t}; F_{1..T}) := \mathbb{E}_{\theta_{t+1..T} \sim F_{t+1..T}}[\sum_{\tau=t+1}^{T} u_\tau(\theta_\tau; \hat{\theta}_{1..t}, \theta_{t+1..\tau}; F_{1..T})].$$

A well-known fact in dynamic mechanism design is that the property DIC implies that the agent's expected overall utility $U_0 = \mathbb{E}[\sum_t u_t(\theta_t; \hat{\theta}_{1..t})]$ is maximized when the agent reports truthfully in each period.

Participation Constraints We will also enforce participation constraints, which require the agent to derive non-negative utility from the mechanism. Again inspired by our main motivation, we will enforce those constraints ex-post, i.e., in every realization of the agent types. It is desirable in an auction setting where the buyers and sellers have an ongoing business relationship (e.g., Internet advertisement) to ensure that buyers derive non-negative utility from the auction. We will refer to this constraint as *ex-post individual rationality*:

$$\sum_{t=1}^{T} u_t(\theta_t; \theta_{1..t}; F_{1..T}) \geq 0. \tag{eP-IR}$$

Revenue optimization We will focus on the problem of maximizing revenue subject to DIC and eP-IR constraints. Fixing a set of distributions $F_{1..T}$, we can define the revenue-optimal mechanism for those distributions as:

$$\textsc{Rev}^*(F_{1..T}) := \max \mathbb{E}_{\theta_{1..T} \sim F_{1..T}}[\sum_{t=1}^{T} p_t(\theta_{1..t}; F_{1..T})] \text{ s.t. (DIC) and (eP-IR).} \tag{\textsc{RevMax}}$$

Static mechanisms Informally, a mechanism is said to be static if the allocation and pricing functions x_t, p_t at time t depend only on the distributional knowledge $F_{1..T}$ and the reported type θ_t at that period. This can be made formal in a measure-theoretic sense by asking x_t and p_t to be measurable with respect to the σ-algebra generated by $(\theta_t, F_{1..T})$.

Under this definition, the revenue optimization problem restricted to static mechanisms becomes separable: the optimal solution consists in applying for each period t the optimal mechanism for that period, i.e., the mechanism $x_t(\theta_t; F_T), p_t(\theta_t; F_t)$ that maximizes $\mathbb{E}_{\theta_t \sim F_t}[p_t(\theta_t, F_t)]$ subject to single-period incentive compatibility and individual rationality.

We can define $\textsc{Rev}^S(F_{1..T})$ as the revenue of the optimal static mechanism. Since the static problem is more constrained, we clearly have $\text{Rev}^*(F_{1..T}) \geq \text{Rev}^S(F_{1..T})$. Papadimitriou et al. [PPPR16] show that the ratio $\text{Rev}^*(F_{1..T})/\text{Rev}^S(F_{1..T})$ can be arbitrarily large. We reproduce this example in Appendix A.

Single item per period setting Although most of our results hold for a general setting, we will present them for single item auctions. In the auction setting the agent type is a single real number representing how much she values one unit of the good, i.e., $\Theta = \mathbb{R}_+$. The outcome $\mathcal{O} = [0, 1]$ corresponds to the probability that the item is allocated to the agent. The value corresponds to a product $v(\theta_t, x_t) = \theta_t \cdot x_t$.

2.1 Non-Clairvoyant Mechanism Design

Informally, a mechanism is *non-clairvoyant* if it depends on distributional knowledge about the present and past, but not the future. In other words, information about the future is not *contractible* (i.e., contract terms explicitly conditional on it).

This notion can be made precise in a measure-theoretic sense: the allocation and pricing function are measurable with respect to the σ-algebra induced by $(\theta_{1..t}, F_{1..t})$. This means that x_t and p_t can be written as: $x_t(\theta_{1..t}; F_{1..t})$ instead of $x_t(\theta_{1..t}; F_{1..T})$.

The notions of DIC and eP-IR are mathematically the same as before. Yet, when DIC and non-clairvoyance are considered together, we obtain a stronger notion of incentive compatibility that allows buyers to verify incentive compatibility without knowledge about the future or agreement between the buyer and seller about future distributions. We illustrate this using an example:

Example 2.1. Consider a setting with a single buyer, two periods and one item being sold per period. The following is a non-clairvoyant incentive compatible mechanism:

- Period 1: elicit type $\hat{\theta}_1$ of the buyer, and give the item for free.
- Period 2: charge $\min(\mathbb{E}_{\theta_2 \sim F_2}[\theta_2], \hat{\theta}_1)$ in advance, and run a second price auction with reserve r such that

$$\mathbb{E}_{\theta_2 \sim F_2}[\max(0, \theta_2 - r)] = \min(\mathbb{E}_{\theta_2 \sim F_2}[\theta_2], \hat{\theta}_1).$$

First, we note that the mechanism is non-clairvoyant since it uses no information about F_2 in the first period. Now notice that fixed any F_1, F_2 we can easily verify that the mechanism satisfies DIC.

Now, let's look at this mechanism from the perspective of the buyer. In period 1, the buyer wants to verify if it is indeed optimal for her to report her type truthfully. She must do so without knowing F_2 since this information is not available to either the buyer or the seller in period 1. *Since reporting truthfully is optimal for any distribution F_2, reporting truthfully is also optimal without knowledge of F_2.*

The last sentence in Example 2.1 hints at a more general phenomenon. Under a non-clairvoyant DIC mechanism, truthtelling is still a best response even if buyers have $\mathcal{F}$-Knightian uncertainty [Kni12] about the future distributions, i.e., if the buyer knows the distribution F_t from which his type at time t is drawn, but for future types at time $t' > t$, the buyer only knows that $F_{t'} \in \mathcal{F}$ for a certain feasible set of distributions $\mathcal{F}$. Formally,

Lemma 2.2. *Let $\mathcal{F}$ be a set of all non-negative measures with a bounded mean. In a DIC non-clairvoyant mechanism, a buyer can implement the optimal strategy by bidding truthfully even with $\mathcal{F}$-Knightian uncertainty about the future.*

Proof. It follows from the same argument as the one in Example 2.1. Given any knowledge about the future, bidding truthfully in the current period is the optimal strategy for the buyer because of the dynamic incentive compatibility. With $\mathcal{F}$-Knightian uncertainty about the future, the set of feasible buyer strategies becomes smaller, while bidding truthfully is still feasible. Therefore, bidding truthfully is still the optimal strategy under $\mathcal{F}$-Knightian uncertainty about the future. □

Non-clairvoyant Revenue Maximization We defined $\text{Rev}^*(F_{1..T})$ as the optimal revenue of a dynamic auction for a sequence of distributions $F_{1..T}$ without imposing any measurability constraints. We call this quantity the *optimal clairvoyant revenue* for $F_{1..T}$. In that section we also defined $\text{Rev}^S(F_{1..T})$ as the revenue of the optimal static mechanism.

The optimal non-clairvoyant revenue for a sequence $F_{1..T}$ is not well-defined, since due to the non-clairvoyance constraint, the incentive constraint is not separable across different distribution sequences. Instead we will define a non-clairvoyant revenue approximation.

Given a certain non-clairvoyant mechanism $\mathcal{M}$, we define its revenue on a sequence of distributions $F_{1..T}$ in the natural way:

$$\textsc{Rev}^{\mathcal{M}}(F_{1..T}) = \mathbb{E}_{\theta_{1..T}\sim F_{1..T}}[\textstyle\sum_t p^{\mathcal{M}}(\theta_{1..t}; F_{1..t})].$$

We say that the non-clairvoyant dynamic mechanism $\mathcal{M}$ is an α-approximation to the clairvoyant benchmark if, for all sequences of distributions $F_{1..T}$,

$$\textsc{Rev}^{\mathcal{M}}(F_{1..T}) \geq \tfrac{1}{\alpha} \cdot \textsc{Rev}^{*}(F_{1..T}).$$

The main question in this paper is whether we can design non-clairvoyant mechanisms that provide good approximations. The optimal static mechanism is non-clairvoyant, but the example in Appendix A shows that it fails to guarantee any approximation α. Given that fact, it is not clear in principle if we can obtain $\alpha < \infty$ at all.

Finally, while there is no notion of the optimal non-clairvoyant mechanism for a sequence $F_{1..T}$, we can define optimality in a maximin sense. Given T and a family of distributions $\mathcal{F}$, we define the maximin optimal non-clairvoyant mechanism as follows (where $0/0 = 1$):

$$\max_{\mathcal{M}} \min_{F_{1..T}\in\mathcal{F}^T} \frac{\textsc{Rev}^{\mathcal{M}}(F_{1..T})}{\textsc{Rev}^{*}(F_{1..T})} \text{ s.t. DIC, eP-IR, } \mathcal{M} \text{ non-clairvoyant} \qquad \text{(MaxiMin)}$$

3 Bank Account Mechanisms

Now we define a general family of auctions called *bank account mechanisms*. We choose this name since they are based on the thought experiment where a buyer "deposits" part of this utility in an account as an investment, which will result in a more favorable auction in future periods. The idea of a bank account is only an abstract device used in the construction of the mechanism and not a real entity that buyers reason about. We initially present our definition in the standard clairvoyant, where there is a fixed sequence of distributions $F_{1..T}$, and allow all functions defining the mechanism to depend on all distributions. To avoid overloading notation, we omit the distribution dependence.

Our auction will have two salient features: (i) each period depends on the previous periods *only* through a single scalar variable called *balance*; and (ii) in this framework, the designer needs to specify single-period auctions that are single-period incentive compatible together with a valid balance update policy. That is, once a valid balance update policy is in place, all the designer needs to worry about are single-period incentive compatibility constraints.

A bank account mechanism B in terms of the following functions for each period:

- A static single-period mechanism $x_t^B(\theta_t, b), p_t^B(\theta_t, b)$ parameterized by a balance $b \in \mathbb{R}_+$ that is (single-period) *incentive-compatible* for each b, i.e.,:

$$v(\theta_t, x_t^B(\theta_t, b)) - p_t^B(\theta_t, b) \geq v(\theta_t, x_t^B(\theta_t', b)) - p_t^B(\theta_t', b). \tag{IC}$$

  Note that we don't require the mechanism to be (single-period) individually rational. We also require the utility of the agent to be *balance independent in expectation*, i.e.,:

$$\mathbb{E}_{\theta_t}[v(\theta_t, x_t^B(\theta_t, b)) - p_t^B(\theta_t, b)] \text{ is a non-negative constant not depending on } b. \tag{BI}$$

- A balance update policy $b_t^B(\theta_t, b)$ which maps the previous balance and the report to the current balance, satisfying the following *balance update* conditions:

$$0 \leq b_t^B(\theta_t, b) \leq b + v(\theta_t, x_t^B(\theta_t, b)) - p_t^B(\theta_t, b). \tag{BU}$$

Given the balance update functions, we can define $b_t : \Theta^t \to \mathbb{R}_+$ recursively as:

$$b_0 = 0 \quad \text{and} \quad b_1(\theta_1) = b_1^B(\theta_1, 0) \quad \text{and} \quad b_t(\theta_{1..t}) = b_t^B(\theta_t, b_{t-1}^B(\theta_{1..t-1})),$$

which allows us to define a dynamic mechanism in the standard sense as:

$$x_t(\theta_{1..t}) = x_t^B(\theta_t, b_{t-1}(\theta_{1..t-1})) \qquad p_t(\theta_{1..t}) = p_t^B(\theta_t, b_{t-1}(\theta_{1..t-1})).$$

In what follows, we will abuse notations by dropping the superscript B and refer to

$x_t(\theta_{1..t})$ and $x_t(\theta_t, b_{t-1})$ interchangeably. Our first theorem is that any bank account mechanism satisfies DIC and eP-IR. In fact, it also satisfied slightly stronger versions of those properties which we discuss in Appendix B.1.

Lemma 3.1. *Any bank account mechanism satisfying IC, BI, and BU is dynamic incentive compatible (DIC) and ex-post individually rational (eP-IR).*

A formal proof of Lemma 3.1 is given in Appendix B.2. Here we provide an intuition of why this mechanism is incentive compatible. We note that when deciding on a strategy in each period, the agent needs to worry about two things: (i) the utility she obtains in this period, which corresponds to $\theta_t x_t - p_t$, and (ii) how her strategy in this period will affect the subsequent periods.

The condition (IC) guarantees that the strategy which maximizes the agent's utility in the current period is to bid her type. The only reason that the agent might think of deviating is to improve her utility in future periods.

The only way that the agent can affect a future period is through the balance. The key idea behind the bank account mechanisms is that the (BI) condition makes the agent do not care (in expectation) about what the balance b_t will be in future periods. Those two facts together guarantee that the mechanism is DIC.

Finally, (eP-IR) follows from summing the condition (BU) over all periods.

The reason we focus on bank account mechanisms and the reason they are useful both in designing optimal dynamic mechanisms and proving lower bounds is that any dynamic incentive compatible and ex-post individually rational mechanism can be converted to a bank account mechanism without loss of revenue or welfare. Therefore, in designing or characterizing the revenue optimal mechanism, it is enough to focus on the subclass of bank account mechanisms. Formally:

Lemma 3.2. *Given any dynamic mechanism $(x_t, p_t)_t$ satisfying DIC and eP-IR, there exists a bank account mechanism with at least the same revenue and at least the same welfare.*

The proof has three main components: (1) we start by transforming a generic DIC and eP-IR mechanism into a mechanism that is still DIC and the agent has zero utility in all but the last period, where her utility is non-negative. We call this a *payment frontloading* mechanism.

(2) The next step is a symmetrization lemma that transforms the mechanism in such a way that if two histories have the same expected total utility in subsequent periods, then their allocation and payments are also the same. (3) The final step shows an isomorphism between payment frontloading symmetric mechanisms and bank account mechanisms. In all those transformations, the revenue and welfare of the mechanism are guaranteed to never decrease.

4 A non-clairvoyant 3-approximation

We describe the central mechanism in the paper for a single buyer case and any number of periods, which is a non-clairvoyant 3-approximation to the revenue of the optimal clairvoyant mechanism. In Section 5 we adapt this mechanism to obtain the maximin optimal mechanism for 2 periods. Later in Section 6 we generalize this result for multiple buyers. Our main result is a non-clairvoyant mechanism that is a 3-approximation to the revenue of the optimal clairvoyant mechanism.

NonClairvoyantBalance Mechanism The mechanism is a combination of three bank account mechanisms. In each period t, it will run a uniform combination of the following three mechanisms:

1. **Give for free:** Allocate the item no matter what the agent type is, and charge her nothing. Increment the balance by her value:

$$x_t^F = 1 \qquad p_t^F = 0 \qquad b_t^F = b_{t-1} + \theta_t.$$

2. **Posted price:** Define a target utility to be $s_t = \min(3b_{t-1}, \mathbb{E}_{\theta_t \sim F_t}[\theta_t])$. Charge this amount from the agent in advance independently of her report, and deduct this amount from the balance. Then, choose a price r_t such that the utility of the agent under r_t is s_t, i.e., $\mathbb{E}_{\theta_t \sim F_t}[(\theta_t - r_t)^+] = s_t$. Since $s_t \le \mathbb{E}[\theta_t]$, the price r_t will be non-negative. Run a posted price auction with this price:

$$x_t^P = \mathbf{1}\{\theta_t \ge r_t\} \qquad p_t^P = s_t + r_t \cdot \mathbf{1}\{\theta_t \ge r_t\} \qquad b_t^P = b_{t-1} - s_t.$$

3. **Myerson's auction:** Find the posted price r_t^* that maximizes the revenue that can be obtained from this period, i.e., $r_t^* = \arg\max_r r \cdot \Pr[\theta_t \geq r]$ and post price r_t^*:

$$x_t^M = \mathbf{1}\{\theta_t \geq r_t^*\} \qquad p_t^M = r_t^* \cdot \mathbf{1}\{\theta_t \geq r_t^*\} \qquad b_t^M = b_{t-1}.$$

We describe the mechanism in each period as a uniform combination of those three:

$$x_t = \tfrac{1}{3}\left[x_t^F + x_t^P + x_t^M\right] \qquad p_t = \tfrac{1}{3}\left[p_t^F + p_t^P + p_t^M\right] \qquad b_t = \tfrac{1}{3}\left[b_t^F + b_t^P + b_t^M\right],$$

where the functions above are functions of $x_t(F_t, \theta_t, b_{t-1}), p_t(F_t, \theta_t, b_{t-1}), b_t(F_t, \theta_t, b_{t-1})$.

The reason why this mechanism is non-clairvoyant is straightforward, since the allocation and payment rule in period t depends only on θ_t, F_t, and the balance b_{t-1} carried from the previous periods, which is itself a function of $\theta_{1..t-1}, F_{1..t-1}$. The mechanism is also ex-post individually rational and dynamic incentive compatible, since it is a bank account mechanism.

Next we show that this mechanism is a 3-approximation. It will be useful to define the notion of the *spend* and use it to show a revenue decomposition lemma. We define the spend as

$$s_t(b_{t-1}) = \max(0, -\min_{\theta_t} v(\theta_t, x_t(\theta_t, b_{t-1})) - p_t(\theta_t, b_{t-1})). \qquad \text{(Spend)}$$

Lemma 4.1 (Revenue upper bound)**.** *The revenue of any dynamic mechanism with a spend function s_t can be bounded by $\mathbb{E}[\sum_t s_t(\theta_{1..t})]$ plus the revenue of the optimal static mechanism.*

Proof. Let $p'_t = p_t - s_t$ we have that: $\text{Rev} \leq \mathbb{E}[\sum_t s_t(\theta_{1..t})] + \mathbb{E}[\sum_t p'_t(\theta_{1..t})]$. Since (x_t, p'_t) is a IC and IR single-period mechanism then its revenue $\mathbb{E}_{\theta_t \sim F_t}[p'_t(\theta_t, b_{t-1})]$ can be bounded by the revenue of the optimal single-period static mechanism for that distribution. □

Theorem 4.2. *The revenue of the* NonClairvoyantBalance *mechanism is at least* 1/3 *of the revenue of the optimal dynamic mechanism.*

Proof. The revenue of the NonClairvoyantBalance mechanism is

$$\textsc{Rev} = \mathbb{E}\left[\sum_t \tfrac{1}{3}[p_t^F(\theta_{1..t}) + p_t^P(\theta_{1..t}) + p_t^M(\theta_{1..t})]\right].$$

Clearly, $\mathbb{E}[\sum_t p_t^M(\theta_{1..t})]$ is the revenue of the optimal static mechanism, in this case the Myerson auction. So by Lemma 4.1, all we need to prove is that $\mathbb{E}[\sum_t p_t^P(\theta_{1..t}) + p_t^F(\theta_{1..t})]$ is greater than the sum of the spends of any optimal bank account mechanism. We will show the stronger statement that for any realization of types $\theta_{1..T}$ we have:

$$\sum_t p_t^P(\theta_{1..t}) + p_t^F(\theta_{1..t}) \geq \sum_t s_t(b_{t-1}(\theta_{1..t-1})) \tag{$*$}$$

Since the realization of the random variables is fixed, let's abbreviate the balance, payment and spend in the generic bank account mechanism by b_t, p_t and s_t. If u_t is the utility of the buyer in period t, define $u'_t = u_t + s_t$. By equations (BU) and (Spend) we know that $0 \leq u'_t \leq \theta_t$, $s_t \leq b_{t-1}$, and $s_t \leq \mathbb{E}_{\hat{\theta}_t \sim F_t}[u'_t(\theta_{1..t-1}, \hat{\theta}_t)] =: \lambda_t$ so:

$$b_t \leq b_{t-1} + u'_t - s_t \qquad u'_t \leq \theta_t \qquad s_t \leq \min(\lambda_t, b_{t-1}) \tag{BalConst}$$

The way to pick u'_t, s_t to optimize $\sum_t s_t$ subject to BalConst is to use the greedy algorithm that always make u'_t as large as possible, i.e., $u'_t = \theta_t$ and always spends as much as possible $s_t = \min(\lambda_t, b_{t-1})$. It should be clear from the principle of local optimality that it is never useful to delay spending outstanding balance. Finally notice that the NonClairvoyant-Balance mechanism implements exactly the optimal Greedy policy scaled by a factor of 1/3: the Give For Free Mechanism adds $\frac{1}{3}\theta_t$ to the balance and the Posted Price Mechanism consumes $\min(b_{t-1}, \frac{1}{3}\lambda_t)$, proving ($*$). Those two facts together prove the theorem. □

5 Maximin optimal mechanisms

Here we derive the maximin optimal mechanism for 2 periods. We start by showing that there is an inherent gap between clairvoyant and non-clairvoyant mechanisms. Formally, we show that no non-clairvoyant mechanism can provide a better-than-2 approximation to the clairvoyant benchmark. Then we show a two-period mechanism that achieves this approximation.

Our lower bound is based on the following idea. Consider a pair of distributions F_1, F_2 and two possible situations: (i) only one item with distribution F_1; and (ii) an item with distribution F_1 followed by another item of distribution F_2. The non-clairvoyant mechanism must allocate the same way in both cases. If the non-clairvoyant mechanism receives a second item, he can allocate and charge a payment for it; however, if not, his revenue will be the one obtained from the first item.

We recall that given a sequence of distributions $F_{1..T}$ we denote by $\textsc{Rev}^*(F_{1..T})$ the revenue of the optimal clairvoyant mechanism and given a non-clairvoyant mechanism M defined by $x_t(\theta_{1..t}; F_{1..t})$ and $p_t(\theta_{1..t}; F_{1..t})$ we define its revenue on a sequence of distributions $F_{1..T}$ by $\textsc{Rev}^M(F_{1..T})$. Given this definitions we prove the following lower bound:

Theorem 5.1 (Lower bound)**.** *For every $\delta > 0$ there are distributions F_1, F_2 such that for every non-clairvoyant mechanism M either $\textsc{Rev}^M(F_1) \leq \frac{1+\delta}{2}\textsc{Rev}^*(F_1)$ or $\textsc{Rev}^M(F_1, F_2) \leq \frac{1+\delta}{2}\textsc{Rev}^*(F_1, F_2)$. In particular, if a non-clairvoyant mechanism is an α-approximation to the clairvoyant benchmark, then $\alpha \geq 2$.*

The central ingredient in the proof (given in Appendix D) is a characterization of non-clairvoyant mechanisms as bank account mechanisms. We define a non-clairvoyant bank account mechanism as a bank account mechanism with the measure-theoretic restriction that the allocation and payment function at time t must be measurable with respect to the balance b_t, the reported type θ_t and the sequence of distributions $F_{1..t}$ corresponding to the current and past periods. In other words, it is simply a bank account mechanism that is not allowed to depend on distributional knowledge about the future.

Our main characterization is that any non-clairvoyant mechanism can be written as a non-clairvoyant bank account mechanism with the same revenue:

Lemma 5.2. *Given any non-clairvoyant dynamic mechanism satisfying DIC and eP-IR, there exists a non-clairvoyant bank account mechanism with the same revenue.*

The characterization in Lemma 5.2 is a non-clairvoyant analogue of Lemma 3.2, and although their proofs share some similarities, there are new challenges to overcome due to the measure-theoretic restrictions imposed by non-clairvoyance: notably the proof of Lemma 3.2 starts by changing the original mechanism to an equivalent payment frontloading mechanism. This clearly breaks non-clairvoyance, so any non-clairvoyant reduction must avoid this step. Also, in the proof of Lemma 3.2, we symmetrize the mechanism around the concept of partially realized utility, which is not well-defined for non-clairvoyant mechanisms. To overcome those problems, we will use two ideas. The first is a strong property implied by non-clairvoyance, which is the fact that the continuation utility must be constant in the reported type (Lemma D.1). The second idea is to symmetrize the mechanism by re-sampling types of previous periods conditioned on a certain event, which in a way resembles the Myersonian ironing procedure.

Maximin optimal mechanism We now present a two-period non-clairvoyant bank account mechanism achieving the optimal approximation for 2 periods. The mechanism uses the same components as the NonClairvoyantBalance mechanism in Section 4, but with different probabilities. In the first period, we run a uniform combination of the *Give for free* and *Myerson* mechanisms:

$$x_1 = \frac{1}{2}[x_1^F + x_1^M] \quad p_1 = \frac{1}{2}[p_1^F + p_1^M] \quad b_1 = \frac{1}{2}[b_1^F + b_1^M],$$

using the notation defined in Section 4. For the second period, we use a uniform combination of *Myerson* and the *Posted Price* auction, with the difference that s_t in the Posted Price mechanism is now defined as $s_2 = \min(2b_1, \mathbb{E}_{\theta_2 \sim F_2}[\theta_2])$. We have

$$x_2 = \frac{1}{2}[x_2^P + x_2^M] \quad p_2 = \frac{1}{2}[p_2^P + p_2^M] \quad b_2 = \frac{1}{2}[b_2^P + b_2^M].$$

Theorem 5.3. *The mechanism described above is the optimal solution to problem (MaxiMin) for $T = 2$ periods.*

Proof. By Theorem 5.1, the solution of the problem (MaxiMin) is at most 1/2, so to show that the mechanism defined above is optimal, we need to argue that it is a 1/2-approximation. This follows from using the same revenue decomposition used in the proof of Theorem 4.2. For two periods, it is easy to write explicitly the revenue upper bound of any (clairvoyant) mechanism in Lemma 4.1: it is at most the static revenue (i.e., $\mathbb{E}[p_1^M(\theta_1)] + \mathbb{E}[p_1^M(\theta_2)]$) plus $\mathbb{E}_{F_1}[\min(\theta_1, \mathbb{E}_{F_2}[\theta_2])]$, which corresponds to the maximum of $\mathbb{E}[s_1 + s_2]$ subject to the constraint (BalConst) in the proof of Theorem 4.2. The non-clairvoyant mechanism described obtains exactly half of that revenue, where the Myerson component obtains half of the optimal static revenue and the *Posted Price* mechanism in the second period obtains at least $\frac{1}{2}\mathbb{E}_{F_1}[\min(\theta_1, \mathbb{E}_{F_2}[\theta_2])]$. □

In Theorem 6.3 we show that it is also possible to obtain a two-period maximin optimal mechanism for any number of buyers.

Maximin optimal forgetful bank account mechanism The previous theorem describes the maximin optimal mechanism among all the non-clairvoyant mechanisms satisfying (DIC) and (eP-IR). Restricting to bank account mechanisms is without loss of generality according Lemma 5.2. The NonClairvoyantBalance mechanism, however, satisfies a more stringent measure-theoretical restriction. While a non-clairvoyant mechanism only requires that the allocation and payment depend only on $\theta_t, b_{t-1}, F_{1..t}$, the NonClairvoyantBalance depends on $F_{1..t-1}$ only through the balance. We call a bank account mechanism whose allocation and payments can be written as a function of (θ_t, b_{t-1}, F_t) a forgetful non-clairvoyant bank account mechanism. This is an especially appealing class of mechanism since it allows the auctioneer to keep track of a single variable, instead of keeping track of a sequence of distributions. The following result, which is proved in Appendix D.3 shows that the NonClairvoyantBalance is the optimal mechanism in this class.

Theorem 5.4. *For $T > 2$, the* NonClairvoyantBalance *is the maximin optimal forgetful bank account mechanism; in other words, it is the solution of the maximin problem* $\max_{\mathcal{M}} \min_{F_{1..T} \in \mathcal{F}^T} \text{Rev}^{\mathcal{M}}(F_{1..T}) / \text{Rev}^*(F_{1..T})$ *subject to the constraint that $\mathcal{M}$ is a forgetful non-clairvoyant bank account mechanism.*

6 Multiple Buyers

In this section, we extend our results to multiple buyer cases. Our decision to focus on a single buyer was driven by the desire to keep notation as simple as possible and to focus on the complications introduced by non-clairvoyance. Once the single buyer case is understood, however, most of the results presented so far extend to the multi-buyer setting. Our characterization results (Lemma 3.2 and Lemma 5.2) extend with essentially no change in the proofs. The lower bound also naturally extends. The only major difference is in the extension of the NonClairvoyantBalance mechanism. Now we need to keep a balance for every buyer, so the state will be a vector. As a consequence, we will be required to reason about utility tradeoffs not only across time periods but across buyers. In the single buyer case, we solved this problem by decreasing the posted price of each buyer based on the bank balance in a greedy manner. Here, instead, we will need to be more careful and decide which auction to use based on the result of an optimization program. This program will resemble what is often called the optimal money burning auction [HR08].

A formal definition of the mechanism design problem for multiple buyers is given in Appendix E.1. It is the natural extension of the single-buyer model with the following incentive notion, which we call Dynamic Bayesian Incentive Compatibility:

$$\theta_t = \arg\max_{\hat{\theta}_t} \mathbb{E}_{\theta_t^{-i}} \left[u_t^i(\theta_t^i; \hat{\theta}_{1..t-1}, (\theta_t^{-i}, \hat{\theta}_t^i)) + U_t^i(\hat{\theta}_{1..t-1}, (\theta_t^{-i}, \hat{\theta}_t^i)) \right], \tag{DBIC}$$

where $U_t^i(\hat{\theta}_{1..t}; F_{1..T})$ is the expected total utility of a buyer in periods $t+1$ to T if her history of reports up to period t is $\hat{\theta}_{1..t}$ and all the buyers report truthfully from period $t+1$ onwards.

The notion of the bank account mechanism can also be naturally extended to multiple buyers. The balance is now an n-dimensional variable $b \in \mathbb{R}_+^n$ and the mechanism in each period is a static IC mechanism parameterized by b satisfying the multi-buyer version of conditions (BI) and (BU). We refer to Appendix E.2 for the details.

6.1 A non-clairvoyant 5-approximation for multiple buyers

Now, we are ready to extend the NonClairvoyantBalance mechanism defined in Section 4 to multiple buyers. We are back to the auction setting where one item is sold per period and the type in each period is a non-negative real number $\theta_t^i \in \mathbb{R}_+$.

We start by observing that Lemma 4.1 still holds in the multi-buyer case. The revenue of any bank account mechanism can be bounded by the revenue of the optimal static mechanism plus the sum of spends $\mathbb{E}[\sum_t \sum_i s_t^i(\theta_{1..t})]$ (see Appendix E.2). A natural strategy given this lemma is to combine the optimal static mechanism (in this case the Myerson auction) with the mechanism that tries to spend as much as possible from the bank accounts. To this end we replace the *give for free* mechanism by a *second price auction* and we replace the *posted price* by the *money burning* mechanism [HR08].

Now, we are ready to define the multi-buyer version of the NonClairvoyantBalance mechanism. As before we will define three mechanisms that are parameterized by the balance b_t together with a balance update policy. As done in Section 4, we will count the spend as part of the payment:

1. **Second Price Auction:** We will allocate the item to the buyer with the highest type (breaking ties arbitrarily). We will increase the balance of the top bidder by her utility. In other words, if we order the buyers such that $\theta_t^1 \geq \theta_t^2 \geq \ldots \geq \theta_t^n$, then:

$$x_t^{S,1} = 1, \quad x_t^{S,j} = 0 \qquad p_t^{S,1} = \theta_t^2, \quad p_t^{S,j} = 0 \qquad b_t^{S,1} = b_{t-1}^1 + \theta_t^1 - \theta_t^2, \quad b_t^{S,j} = b_{t-1}^j$$

   for $j \geq 2$. The mechanism guarantees the largest possible increase in bank balance.

2. **Money Burning Auction:** Given the bank account states b_{t-1} we will compute the single-period mechanism that maximizes the sum of expected utilities of the buyers subject to each buyer i having utility at most $\frac{5}{2}b_{t-1}^i$, this is, we want to compute the allocation and payment rule $x_t^B, \tilde{p}_t^B$ satisfying Bayesian incentive compatibility and individual rationality and maximizing:

$$\max \textstyle\sum_i \mathbb{E}_{\theta_t}[\tilde{u}_t^{B,i}(\theta_t)] \quad \text{s.t.} \quad \mathbb{E}[\tilde{u}_t^{B,i}] \leq \frac{5}{2}b_{t-1}^i, \forall i \qquad \text{(BIC) and (IR)}$$

   Money burning mechanisms have this name since they correspond to the welfare maximization problem when the revenue obtained is burned. Hartline and Roughgarden [HR08] provide a comprehensive study of such mechanisms and show that they can be written as a virtual value maximization for a different notion of virtual values. In fact we can deduce from their result that the solution to the problem above corresponds to the auction where we transform the values to the space of virtual values for utilities and run a (scaled) second price auction in that space. So in that sense it is not very different from Myerson's auction other than the fact that the notion of virtual values is non-standard. Given such a solution, we define the money burning mechanism using the allocation obtained from the program and payment and balance as follows:

$$p_t^{B,i} = \tilde{p}_t^{B,i} + \mathbb{E}[\tilde{u}_t^{B,i}] \qquad b_t^{B,i} = b_{t-1}^i - \mathbb{E}[\tilde{u}_t^{B,i}]$$

3. **Myerson's Auction:** We run the static optimal auction given by x^M and p^M. Bank accounts are unchanged, i.e., $b_t^{M,i} = b_{t-1}^i$.

Now, the non-clairvoyant balance mechanism is the mechanism defined by:

$$x_t^i = \tfrac{1}{5}x_t^{M,i} + \tfrac{2}{5}x_t^{S,i} + \tfrac{2}{5}x_t^{B,i} \qquad p_t^i = \tfrac{1}{5}p_t^{M,i} + \tfrac{2}{5}p_t^{S,i} + \tfrac{2}{5}p_t^{B,i} \qquad b_t^i = \tfrac{1}{5}b_t^{M,i} + \tfrac{2}{5}b_t^{S,i} + \tfrac{2}{5}b_t^{B,i}$$

In Appendix F we argue that each component of the NonClairvoyantBalance mechanism can be implemented as a virtual value maximizer. Next we provide an approximation guarantee with respect to the clairvoyant benchmark:

Theorem 6.1. *The multi-buyer version of the* NonClairvoyantBalance *mechanism is a non-clairvoyant* 5*-approximation to the clairvoyant benchmark.*

Stronger incentive guarantees. While our mechanism provides 1/5 of the revenue of any Dynamic Bayesian Incentive Compatible (DBIC) mechanism, it actually satisfies a stronger notion of incentive compatibility: it is optimal for an agent to report her true type even if she knows the types of other agents in the period when she is reporting. This corresponds to the notion of Strong Dynamic Bayesian Incentive Compatibility:

$$\theta_t = \arg\max_{\hat{\theta}_t} u_t^i(\theta_t^i; \hat{\theta}_{1..t-1}, (\theta_t^{-i}, \hat{\theta}_t^i)) + U_t^i(\hat{\theta}_{1..t-1}, (\theta_t^{-i}, \hat{\theta}_t^i)), \quad \forall \theta_t^{-i} \tag{sDBIC}$$

Lemma 6.2. *The* NonClairvoyantBalance *mechanism satisfies (sDBIC).*

The proof is straightforward, but we include it in Appendix E.4 for completeness.

6.2 Maximin optimal multi-buyer mechanism for two periods

Finally we adapt the 2-period mechanism in Section 5 to multiple buyers and claim it is still a 2-approximation and therefore it is maximin optimal. The mechanism has the same three components we used for the multi-buyer version of NonClairvoyantBalance, except that the Money Burning Auction is slightly different on the coefficients (2 instead of 5/2) in the spend constraints:

$$\max \textstyle\sum_i \mathbb{E}_{\theta_t}[\tilde{u}_t^B(\theta_t)] \quad \text{s.t.} \quad \mathbb{E}[\tilde{u}_t^B] \leq 2b_{t-1}^i, \forall i \quad \text{(IC) and (IR)}$$

Then the 2-period version of the NonClairvoyantBalance mechanism is defined by:

$$\begin{array}{lll} x_1^i = \frac{1}{2}\left[x_1^{M,i} + x_1^{S,i}\right] & p_1^i = \frac{1}{2}\left[p_1^{M,i} + p_1^{S,i}\right] & b_1^i = \frac{1}{2}\left[b_1^{M,i} + b_1^{S,i}\right] \\ x_2^i = \frac{1}{2}\left[x_2^{M,i} + x_2^{B,i}\right] & p_2^i = \frac{1}{2}\left[p_2^{M,i} + p_2^{B,i}\right] & b_2^i = \frac{1}{2}\left[b_2^{M,i} + b_2^{B,i}\right] \end{array}$$

Theorem 6.3. *The* 2*-period version of the* NonClairvoyantBalance *mechanism is a non-clairvoyant* 2*-approximation to the* 2*-period clairvoyant benchmark. Hence it is an optimal solution to problem (MaxiMin) for* $T = 2$ *periods.*

A proof is included in Appendix E.4.

7 Related Work

Dynamic mechanism design The literature on dynamic mechanism design is too extensive to survey here: we refer to the survey by Bergemann and Said [BS11] and another recent survey by Bergemann and Välimäki [BV17] for comprehensive treatments on the subject. Here, we discuss a few representative papers in the literature.

For efficiency (social-welfare) maximization, Bergemann and Välimäki [BV10] propose the dynamic pivot mechanism, which is a natural generalization of the VCG mechanism to a dynamic environment where agents receive private information over time, and Athey and Segal [AS13] propose the team mechanism to achieve budget-balanced outcomes (see also Bergemann and Välimäki [BV03, BV06], Cavallo, Parkes, and Singh [CPS06, CPS09], and

Cavallo [Cav08]).

For revenue maximization, a line of research was initiated by Baron and Besanko [BB84] and Courty and Li [CH00] that studies the setting where the private information of agents varies over time. The latter show an optimal dynamic contract that "screens" the agents twice in a setting where agents initially have private information about the future distribution of their values (see also [BS12, AAD15] for "screening" in dynamic mechanism design).

Eső and Szentes [ES07] study a closely related two-period model, where the agents only have a rough estimation of their private values to the item in the first round and the seller can release additional signals to affect their values before selling the item in the second round. In a particular setting, they propose a "handicap auction" that shares some similar ideas with our bank account mechanism in each period: in a "handicap auction", the agents buy their premiums from a menu offered by the seller in the first round based on their rough estimation of private values, and then compete with each other under unequal conditions (premiums) in the second round after receiving additional signals from the seller. It is similar to our bank account mechanisms in the sense that in both settings, the agents first buy some advantages/discounts for the next round via either premium costs (in "handicap auctions") or spends (in bank account auctions) based on rough estimations of their values (prior distributions of each period in our case), and then compete under different levels of advantage after observing their realized values.

Pavan, Segal, and Toikka [PST09, PST10, PST14] generalize the idea of Myerson [Mye81] to a multi-period setting with dynamic private information and characterize the incentive compatibility in terms of necessary conditions and some sufficient conditions. Kakade, Lobel, and Nazerzadeh [KLN13] propose the virtual-pivot mechanism by combining ideas of "virtual values" for static optimal mechanism design [Mye81] and "dynamic pivot mechanisms" for dynamic efficient mechanism design [BV10]. In particular, they show that the virtual-pivot mechanisms are optimal in certain dynamic environments that are "separable", satisfy periodic ex-post incentive compatible and individually rational, and have simple structure in multi-armed bandit settings (see also [Bat05, Deb08] for settings with private values evolving through Markovian processes). Devanur, Peres, and Sivan [DPS15] and Chawla, Devanur, Karlin, and Sivan [CDKS16] study the repeated selling of fresh copies of an item to a single buyer who has either fixed private value [DPS15] or evolving values [CDKS16] to the copies.

One major difference between our setting and the one with dynamic private information we just discussed above and is that we have no initial private types for the agents and the private types/values are independent of previous outcomes. Instead, we are able to guarantee *ex-post individual rationality* for a very general setting in our case, while weaker notions of individual rationality (i.e., interim individual rationality or individually rational in expectation) are adopted in most of the previous studies (except for [KLN13], which guarantees ex-post individual rationality for environments satisfying a separability condition).

There are more works primarily focused on the setting with dynamic populations and fixed information [PS03, Gal06, Boa08, PV08, Sai08, GM09, BS10, GM10, Sai12]. In particular, the notion of non-clairvoyance we introduced is similar in spirit with the online mechanism design setting studied by Parkes and Singh [PS03] (for welfare-maximization) and Pai and Vohra [PV08] (for revenue-maximization) in the sense that the designer has restricted information about dynamic arrival/departure (for online mechanisms) or dynamic prior distributions (for non-clairvoyant mechanisms) in future periods. In contrast to the settings with dynamic populations discussed above, however, our setting emphasizes the dynamic arrivals of perishable goods (e.g., ad impressions), while it is still general enough to capture the dynamic attendance of agents by setting periodic prior distributions to be $\Pr[v = 0] = 1$ when they are absent from the auction except that the agents have unlimited demands. Hock [Hoc03] studies the revenue-maximization problem for selling homogeneous items to unit demand buyers where the demand curve is unknown. In particular, he considers an approach of selling the items sequentially and setting the optimal price for the current buyer based on the demand curve estimated from bids of previous buyers, which is also related to our notion of non-clairvoyance.

Our work is closer to the line of inquiry initiated by Papadimitriou et al. [PPPR16], who seek to design revenue-optimal auctions in the setting where items are sequentially sold to the same set of buyers over time. They first show that the problem of designing the optimal deterministic auction is NP-hard even for 1 agent and 2 periods, but they provide a polynomial time algorithm for the optimal randomized auction via a linear programming formulation for a constant number of buyers and correlated valuations. The formulation is exponential in the number of buyers and the support of the distribution of agent type profiles over time. If agents have independent types over periods this causes their formulation to become exponential in the number of periods as well. This problem was addressed by Ashlagi, Daskalakis, and Haghpanah [ADH16], who replaced the linear programming formulation with a dynamic program and obtained a $(1+\epsilon)$-approximation that is polynomial in the number of periods for a single buyer with independent valuations. For multiple buyers they provide a mathematical characterization but not an algorithm to solve it. Simultaneously and independently, we also provide a $(1+\epsilon)$-approximation for agents with independent valuations using dynamic programming in the unpublished manuscript [MLTZ16].

Another closely related stream of literature is the design of dynamic mechanism in the time-discounted model where valuations of the buyers are drawn from an identical distribution in each step. This line was initiated by Biais et al. [BMPR07] and Krishna et al. [KLT13]. Belloni, Chen, and Sun [BCS] provide a characterization of the optimal mechanism by extending Myerson's ironing technique to dynamic settings. Balseiro, Mirrokni, and Paes Leme [BMPL16] study the effect of imposing stronger constraints on the utilities of buyers, and design closed-form mechanisms that approach the optimal in the limit. This line of literature is incomparable with our work: their settings are i.i.d. across time (while we only assume independence), focus on a single buyer, and are based on a fixed-point formulation that is only possible in time-discounted models. While their model is more restricted, they are able to provide stronger guarantees and closed-form mechanisms.

Dynamic mechanism design frameworks One major contribution of our paper is the bank account framework, which provides a general framework to design (traditional or non-clairvoyant) dynamic mechanisms. In particular, incorporating this framework with ex-post individual rationality is technically challenging. Another major step in the development of the bank account framework is to show that all non-clairvoyant mechanisms can be cast in it. There have been other very interesting and useful frameworks, the oldest of which seems to be the *promised utility* framework of Thomas and Worrall [TW90] (see Belloni et al. [BCS16] or Balseiro et al. [BMPL16] for recent applications). More recently, Ashlagi et al. [ADH16] designed a framework based on revenue-utility tradeoff functions. Both the results in [TW90] and [ADH16] accommodate ex-post individual rationality and are universal in the sense that the optimal mechanism is always contained in their class.

The main difference between bank accounts and promised utilities or revenue-utility tradeoffs is that while the latter two are forward-looking (i.e., they define an optimal form for one period, given the optimal solution for the next), the bank account framework is backward-looking. It defines an allocation and pricing rule based on the past and not the future. To the best of our knowledge, this is the only framework capable of accommodating non-clairvoyance.

Online supply and scheduling The term *clairvoyant* is borrowed from the scheduling literature, where it is typically used to refer to an algorithm that can 'see the future' in the sense that it can know, for example, the total execution time of jobs not yet completed. It is also often used to describe an adversary that can predict all the algorithm actions, present and future. The concept of non-clairvoyance is typically used to refer to an algorithm that can perform a certain task well, regardless of having all information.

In that sense, one can see our paper as an *online algorithm approach* to dynamic mechanism design. The study of incentives in problems where items arrive over time in an online manner was initiated by Babaioff, Blumrosen and Roth [BBR10], who design auctions (and prove lower bounds) for problems where incentives are required to be maintained and we are required to allocate goods without information about what the total supply is. This was extended by Goel et al. [GML13] to budgeted settings. The online supply problem was also studied from the perspective of revenue in both the Bayesian and prior free settings by Mahdian and Saberi [MS06] and Devanur and Hartline [DH09]. In this line of work, however, agents review their types in the beginning of the period, and the challenge is to guarantee a monotone allocation. Since types are only reported once, incentive constraints don't need to be enforced dynamically.

Robustness and Detail Independence Non-clairvoyance can be seen as a form of robustification of dynamic mechanisms. By requiring the mechanism not to use any distributional information from future periods, we obtain mechanisms that are much less detail-dependent, in the spirit of Wilson's doctrine [Wil87]. In this sense, we share the philosophy of Bergemann and Morris [BM12] in their theory of robust dynamic mechanism design, which seeks to design mechanisms that work irrespective of beliefs that agents might have. While we make the mechanisms free of beliefs about the future, we still assume beliefs about the present (i.e., the seller in period t has forecast F_t for demand during that period). In that sense we are more in line with Yogi Berra, who says "*It's tough to make predictions, especially about the future.*"

Simpler Mechanisms without Backward Induction The constraints imposed by non-clairvoyance naturally produce simpler mechanisms. To illustrate the simplicity of the NonClairvoyantBalance mechanism, it is useful to compare it with previous approaches to designing dynamic mechanisms. All previous approaches require some form of expensive pre-processing step. In [PPPR16], the allocation and pricing are determined by the solutions of a large linear program that has one variable for each sequence of reports. If the distributions are independent, this requires a number of variables that are exponential both in the number of buyers and the number of periods. Another approach is to replace the linear program by a dynamic program that is solved via backward induction. This is the approach taken by Ashlagi et al [ADH16] and by [MLTZ16]. The mechanism extracts a $(1 - \epsilon)$ fraction of the optimal revenue, but it is no longer exponential in the number of periods. In both cases, it is only analyzed for a single buyer. The mathematical characterization of the optimal mechanism for multiple buyers is also presented in [ADH16], but it is not made algorithmic beyond a single buyer. Ashlagi et al. [ADH16] also propose a second mechanism which extracts at least 1/2 of the optimal revenue but requires solving a simpler dynamic program and produces a simpler allocation rule; however, it still requires backward induction and only applies to one buyer.

Non-clairvoyance clearly prevents the designer from using any form of backward induction, since, at period t, we don't know the distributions in future periods. In fact, we don't even know how many more items we will have to allocate. The NonClairvoyantBalance mechanism requires no backward induction: in each period t, it uses the distributions of the buyers at that period to construct an optimal auction (which is based on virtual values, following the Myersonian approach), a second price auction, and a money burning auction (which also admits a virtual value description).

In summary, we get an auction that requires no pre-processing and no backward induction. Moreover, each of its components is a virtual value maximizer.

References

[AAD15] Mustafa Akan, Barış Ata, and James D Dana. Revenue management by sequential screening. *Journal of Economic Theory*, 159:728–774, 2015. 24

[ADH16] Itai Ashlagi, Constantinos Daskalakis, and Nima Haghpanah. Sequential mechanisms with ex-post participation guarantees. In *EC '16*, pages 213–214, 2016. 26, 28

[AS13] Susan Athey and Ilya Segal. An efficient dynamic mechanism. *Econometrica*, 81(6):2463–2485, 2013. 23, 53

[Bat05] Marco Battaglini. Long-term contracting with markovian consumers. *The American economic review*, 95(3):637–658, 2005. 24

[BB84] David P Baron and David Besanko. Regulation and information in a continuing relationship. *Information Economics and policy*, 1(3):267–302, 1984. 24

[BBR10] Moshe Babaioff, Liad Blumrosen, and Aaron Roth. Auctions with online supply. In *Proceedings of the 11th ACM conference on Electronic commerce*, pages 13–22. ACM, 2010. 27

[BCS] Alexandre Belloni, Bingyao Chen, and Peng Sun. Structures of optimal contracts in dynamic mechanism design with one agent. 26

[BCS16] Alexandre Belloni, Bingyao Chen, and Peng Sun. Computation of optimal dynamic mechanism with participation requirements. Technical report, working paper, 2016. 26

[Ber00] Dimitri P. Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, 2nd edition, 2000. 36

[BM05] Dirk Bergemann and Stephen Morris. Robust mechanism design. *Econometrica*, 73(6):1771–1813, 2005. 53

[BM12] Dirk Bergemann and Stephen Morris. *Robust mechanism design: The role of private information and higher order beliefs*, volume 2. World Scientific, 2012. 27

[BMPL16] Santiago Balseiro, Vahab Mirrokni, and Renato Paes Leme. Dynamic mechanisms with martingale utilities. 2016. 26

[BMPR07] Bruno Biais, Thomas Mariotti, Guillaume Plantin, and Jean-Charles Rochet. Dynamic security design: Convergence to continuous time and asset pricing implications. *The Review of Economic Studies*, 74(2):345–390, 2007. 26

[Boa08] Simon Board. Durable-goods monopoly with varying demand. *The Review of Economic Studies*, 75(2):391–413, 2008. 25

[BS10] Simon Board and Andrzej Skrzypacz. Optimal dynamic auctions for durable goods: Posted prices and fire-sales. *Unpublished manuscript, Stanford University*, 2010. 25

[BS11] Dirk Bergemann and Maher Said. Dynamic auctions. *Wiley Encyclopedia of Operations Research and Management Science*, 2011. 23

[BS12] Raphael Boleslavsky and Maher Said. Progressive screening: Long-term contracting with a privately known stochastic process. *The Review of Economic Studies*, page rds021, 2012. 24

[BV03] Dirk Bergemann and Juuso Välimäki. Dynamic common agency. *Journal of Economic Theory*, 111(1):23–48, 2003. 23

[BV06] Dirk Bergemann and Juuso Välimäki. Dynamic price competition. *Journal of Economic Theory*, 127(1):232–263, 2006. 23

[BV10] Dirk Bergemann and Juuso Välimäki. The dynamic pivot mechanism. *Econometrica*, 78(2):771–789, 2010. 23, 24

[BV17] Dirk Bergemann and Juuso Välimäki. Dynamic mechanism design: An introduction. Technical report, Cowles Foundation for Research in Economics, Yale University, 2017. 23

[Cav08] Ruggiero Cavallo. Efficiency and redistribution in dynamic mechanism design. In *Proceedings of the 9th ACM conference on Electronic commerce*, pages 220–229. ACM, 2008. 24

[CDKS16] Shuchi Chawla, Nikhil R Devanur, Anna R Karlin, and Balasubranianian Sivan. Simple pricing schemes for consumers with evolving values. In *SODA 2016*, pages 1476–1490. SIAM, 2016. 24, 25

[CDW12a] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. An algorithmic characterization of multi-dimensional mechanisms. In *STOC 2012*, pages 459–478, 2012. 58

[CDW12b] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In *FOCS 2012*, 2012. 58

[CH00] Pascal Courty and Li Hao. Sequential screening. *The Review of Economic Studies*, 67(4):697–717, 2000. 24

[CPS06] Ruggiero Cavallo, David C Parkes, and Satinder Singh. Optimal coordinated planning amongst self-interested agents with private state. In *UAI 2006*, pages 55–62. AUAI Press, 2006. 23

[CPS09] Ruggiero Cavallo, David C Parkes, and Satinder Singh. Efficient mechanisms with dynamic populations and dynamic types. 2009. 23

[Deb08] Rahul Deb. Optimal contracting of new experience goods. 2008. 24

[DH09] Nikhil R Devanur and Jason D Hartline. Limited and online supply and the bayesian foundations of prior-free mechanism design. In *EC '09*, pages 41–50. ACM, 2009. 27

[DPS15] Nikhil R Devanur, Yuval Peres, and Balasubramanian Sivan. Perfect bayesian equilibria in repeated sales. In *SODA 2015*, pages 983–1002. SIAM, 2015. 24, 25

[Elk07] Edith Elkind. Designing and learning optimal finite support auctions. In *SODA 2007*, pages 736–745. SIAM, 2007. 59

[ES07] Péter Eső and Balazs Szentes. Optimal information disclosure in auctions and the handicap auction. *The Review of Economic Studies*, 74(3):705–731, 2007. 24

[Gal06] Jérémie Gallien. Dynamic mechanism design for online commerce. *Operations Research*, 54(2):291–310, 2006. 25

[GM09] Alex Gershkov and Benny Moldovanu. Dynamic revenue maximization with heterogeneous objects: A mechanism design approach. *American economic Journal: microeconomics*, 1(2):168–198, 2009. 25

[GM10] Alex Gershkov and Benny Moldovanu. Efficient sequential assignment with incomplete information. *Games and Economic Behavior*, 68(1):144–154, 2010. 25

[GML13] Gagan Goel, Vahab S. Mirrokni, and Renato Paes Leme. Clinching auction with online supply. In *SODA 2013*, pages 605–619, 2013. 27

[Hoc03] Dee Hock. Optimal pricing mechanisms with unknown demand. *The American economic review*, 93(3):509–529, 2003. 25

[HR08] Jason D. Hartline and Tim Roughgarden. Optimal mechanism design and money burning. In *STOC 2008*, pages 75–84, 2008. 5, 20, 21, 58, 59

[JS07] Matthew O Jackson and Hugo F Sonnenschein. Overcoming incentive constraints by linking decisions. *Econometrica*, 75(1):241–257, 2007. 1

[KLN13] Sham M Kakade, Ilan Lobel, and Hamid Nazerzadeh. Optimal dynamic mechanism design and the virtual-pivot mechanism. *Operations Research*, 61(4):837–854, 2013. 24, 25

[KLT13] R Vijay Krishna, Giuseppe Lopomo, and Curtis R Taylor. Stairway to heaven or highway to hell: Liquidity, sweat equity, and the uncertain path to ownership. *The RAND Journal of Economics*, 44(1):104–127, 2013. 26

[Kni12] Frank H Knight. *Risk, uncertainty and profit.* Courier Corporation, 2012. 10

[MLTZ16] Vahab S. Mirrokni, Renato Paes Leme, Pingzhong Tang, and Song Zuo. Optimal dynamic mechanisms with ex-post IR via bank accounts. *CoRR*, abs/1605.08840, 2016. 26, 28

[MR92] Dilip Mookherjee and Stefan Reichelstein. Dominant strategy implementation of bayesian incentive compatible allocation rules. *Journal of Economic Theory*, 56(2):378–399, 1992. 53

[MS06] Mohammad Mahdian and Amin Saberi. Multi-unit auctions with unknown supply. In *Proceedings of the 7th ACM conference on Electronic commerce*, pages 243–249. ACM, 2006. 27

[MV07] Alejandro M Manelli and Daniel R Vincent. Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *Journal of Economic theory*, 137(1):153–185, 2007. 1

[Mye81] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981. 24, 49

[PPPR16] Christos Papadimitriou, George Pierrakos, Christos-Alexandros Psomas, and Aviad Rubinstein. On the complexity of dynamic mechanism design. In *SODA 2016*. SIAM, 2016. 1, 9, 25, 28, 33

[PS03] David C Parkes and Satinder Singh. An mdp-based approach to online mechanism design. In *NIPS 2003*, pages 791–798. MIT Press, 2003. 25

[PST09] Alessandro Pavan, Ilya R Segal, and Juuso Toikka. Dynamic mechanism design: Incentive compatibility, profit maximization and information disclosure. 2009. 24

[PST10] Alessandro Pavan, Ilya Segal, and Juuso Toikka. Infinite-horizon mechanism design: The independent-shock approach. Technical report, Discussion Paper, Center for Mathematical Studies in Economics and Management Science, 2010. 24

[PST14] Alessandro Pavan, Ilya Segal, and Juuso Toikka. Dynamic mechanism design: A myersonian approach. *Econometrica*, 82(2):601–653, 2014. 24, 53

[PV08] Mallesh Pai and Rakesh V Vohra. Optimal dynamic auctions. Technical report, Discussion paper//Center for Mathematical Studies in Economics and Management Science, 2008. 25

[Sai08] Maher Said. Information revelation and random entry in sequential ascending auctions. In *Proceedings of the 9th ACM conference on Electronic commerce*, pages 98–98. ACM, 2008. 25

[Sai12] Maher Said. Auctions with dynamic populations: Efficiency and revenue maximization. *Journal of Economic Theory*, 147(6):2419–2438, 2012. 25

[TW90] Jonathan Thomas and Tim Worrall. Income fluctuation and asymmetric information: An example of a repeated principal-agent problem. *Journal of Economic Theory*, 51(2):367 – 390, 1990. 26, 36

[Wil87] Robert Wilson. Game theoretic analyses of trading processes. In *Advances in Economic Theory: Fifth World Congress, ed. Truman Bewley*. Cambridge: Cambridge University Press, 1987. 27

A Example from Section 2

The following example is essentially the same as the one given by [PPPR16], which shows that the gap between dynamic and static mechanisms could be arbitrarily large.

Example A.1. Consider a clairvoyant dynamic mechanism with two periods and a single item auctioned in each period. In each period, the valuation of the buyer for the item is independently drawn from the *equal revenue distribution*:

$$F_1 = F_2 = F,\ F(\theta) = \begin{cases} 0, & \theta \leq 1 \\ 1 - 1/\theta, & 1 < \theta < \theta_{\max} \\ 1, & \theta \geq \theta_{\max} \end{cases}.$$

Note that the revenue of the optimal static mechanism is 2, which is the revenue obtained by running Myerson's auction in each period. Using a dynamic mechanism we can obtain revenue $2 + \ln\ln\theta_{\max}$. The gap is arbitrarily large as $\theta_{\max}$ tends to infinity.

The dynamic mechanism works as follows:

- In the first period, the seller allocates the item with probability 1 to the buyer and charges her bid but no more than $1 + \ln\theta_{\max}$, i.e., $p_1(\hat{\theta}_1) = \min(\hat{\theta}_1, 1 + \ln\theta_{\max})$.
- In the second period, the seller runs a posted-price mechanism with price $r_2 = \theta_{\max}/e^{p_1 - 1}$.

In fact, the dynamic mechanism above is dynamic incentive compatible and ex-post individually rational. This is because it can be written as a bank account mechanism (that we discuss in Section 3). The expected revenue is,

$$\begin{aligned} \textsc{Rev} &= \mathbb{E}[p_1(\theta_1) + p_2(\theta_1, \theta_2)] = \mathbb{E}[\min(\theta_1, 1 + \ln\theta_{\max})] + \mathbb{E}[r_2 \cdot \mathbf{1}\{\theta_2 \geq r_2\}] \\ &= 2 + \mathbb{E}[\min(\theta_1 - 1, \ln\theta_{\max})] = 2 + \ln\ln\theta_{\max}. \end{aligned}$$

B Proofs from Section 3

B.1 Stronger properties of bank account mechanisms

Besides DIC and eP-IR, bank account mechanisms satisfy stronger versions of both IC and IR, which we describe below. The mechanism is per-period incentive compatible, i.e., the buyer's utility in each given period is maximized by reporting truthfully in that period:

$$\theta_t \in \arg\max_{\hat{\theta}_t} u_t(\theta_t; \theta_{1..t-1}, \hat{\theta}_t), \tag{pp-IC}$$

and the expected continuation utility is independent of the type reported, i.e.:

$$U_t(\theta_{1..t-1}, \hat{\theta}_t) \text{ is independent of } \hat{\theta}_t. \tag{indCont}$$

It is straightforward from the definition of (DIC) that (pp-IC) and (indCont) imply (DIC).

The mechanism also satisfies a stronger version of (eP-IR): it is ex-post individually rational for every prefix and every realization of the random variables:

$$\textstyle\sum_{\tau=1}^{t} u_\tau(\theta_\tau; \theta_{1..\tau}) \geq 0, \forall t. \tag{prefix-epIR}$$

Moreover, each individual period is individually rational in expectation:

$$\mathbb{E}_{\theta_t}[u_t(\theta_t; \theta_{1..t})] \geq 0, \forall t. \tag{$\mathbb{E}$pp-IR}$$

The fact that bank account mechanisms also satisfy (pp-IC), (indCont), (prefix-epIR) and ($\mathbb{E}$pp-IR) follows directly from the proof of Lemma 3.1.

B.2 Proof of Lemma 3.1

Proof of Lemma 3.1. First we prove that conditions (IC) and (BI) imply that the mechanism satisfies (DIC). By definition,

$$\begin{aligned} u_t(\theta_t; \hat{\theta}_{1..t-1}, \hat{\theta}_t) = & \ v(\theta_t, x_t(\hat{\theta}_{1..t-1}, \hat{\theta}_t)) - p_t(\hat{\theta}_{1..t-1}, \hat{\theta}_t) \\ = & \ v(\theta_t, x_t^B(\hat{\theta}_t, b_{t-1}(\hat{\theta}_{1..t-1}))) - p_t^B(\hat{\theta}_t, b_{t-1}(\hat{\theta}_{1..t-1})). \end{aligned}$$

Combined with (IC), we have

$$\begin{aligned} u_t(\theta_t; \hat{\theta}_{1..t-1}, \hat{\theta}_t) = & \ v(\theta_t, x_t^B(\hat{\theta}_t, b_{t-1}(\hat{\theta}_{1..t-1}))) - p_t^B(\hat{\theta}_t, b_{t-1}(\hat{\theta}_{1..t-1})) \\ \leq & \ v(\theta_t, x_t^B(\theta_t, b_{t-1}(\hat{\theta}_{1..t-1}))) - p_t^B(\theta_t, b_{t-1}(\hat{\theta}_{1..t-1})) = u_t(\theta_t; \hat{\theta}_{1..t-1}, \theta_t). \end{aligned} \tag{B.1}$$

By (BI), $\mathbb{E}_{\theta_\tau}[u_\tau(\theta_\tau; \hat{\theta}_{1..t-1}, \hat{\theta}_t, \theta_{t+1..\tau})]$ is constant in $b_{\tau-1} = b_{\tau-1}(\hat{\theta}_{1..t-1}, \hat{\theta}_t, \theta_{t+1..\tau-1})$, hence also constant in $\hat{\theta}_t$, namely,

$$\mathbb{E}_{\theta_\tau}\left[u_\tau(\theta_\tau; \hat{\theta}_{1..t-1}, \hat{\theta}_t, \theta_{t+1..\tau})\right] = \mathbb{E}_{\theta_\tau}\left[u_\tau(\theta_\tau; \hat{\theta}_{1..t-1}, \theta_t, \theta_{t+1..\tau})\right].$$

Therefore

$$\mathbb{E}_{\theta_{t+1..T}}\left[\textstyle\sum_{\tau=t+1}^{T} u_\tau(\theta_\tau; \hat{\theta}_{1..t-1}, \hat{\theta}_t, \theta_{t+1..\tau})\right] = \mathbb{E}_{\theta_{t+1..T}}\left[\textstyle\sum_{\tau=t+1}^{T} u_\tau(\theta_\tau; \hat{\theta}_{1..t-1}, \theta_t, \theta_{t+1..\tau})\right]. \tag{B.2}$$

Adding (B.1) and (B.2) together, we have

$$
\begin{aligned}
& u_t(\theta_t;\hat{\theta}_{1..t-1},\hat{\theta}_t) + \mathbb{E}_{\theta_{t+1..T}}\left[\sum_{\tau=t+1}^{T} u_\tau(\theta_\tau;\hat{\theta}_{1..t-1},\hat{\theta}_t,\theta_{t+1..\tau})\right] \\
\leq\ & u_t(\theta_t;\hat{\theta}_{1..t-1},\theta_t) + \mathbb{E}_{\theta_{t+1..T}}\left[\sum_{\tau=t+1}^{T} u_\tau(\theta_\tau;\hat{\theta}_{1..t-1},\theta_t,\theta_{t+1..\tau})\right],
\end{aligned}
$$

which directly implies (DIC).

Now we show that (BU) implies (eP-IR). Summing up (BU) for $t = 1$ to T, we have,

$$
\begin{aligned}
& \sum_{t=1}^{T} b_t \leq \sum_{t=1}^{T}\left(b_{t-1} + v(\theta_t, x_t^B(\theta_t, b_{t-1})) - p_t^B(\theta_t, b_{t-1})\right), \\
\Longrightarrow\ & \sum_{t=1}^{T}\left(v(\theta_t, x_t^B(\theta_t, b_{t-1})) - p_t^B(\theta_t, b_{t-1})\right) \geq b_T - b_0 \geq 0. \qquad \text{(B.3)}
\end{aligned}
$$

Again, by definition,

$$
u_t(\theta_t;\theta_{1..t}) = v(\theta_t, x_t(\theta_{1..t})) - p_t(\theta_{1..t}) = v(\theta_t, x_t^B(\theta_t, b_{t-1})) - p_t^B(\theta_t, b_{t-1}).
$$

(eP-IR) is then implied by (B.3), i.e.,

$$
\sum_{t=1}^{T} u_t(\theta_t;\theta_{1..t}) = \sum_{t=1}^{T}\left(v(\theta_t, x_t^B(\theta_t, b_{t-1})) - p_t^B(\theta_t, b_{t-1})\right) \geq 0.
$$

□

B.3 Proof of Lemma 3.2

The first step of the proof is a symmetrization lemma. Central to this lemma is the concept of partially-realized utility, which measures the expected utility of an agent conditioned on some prefix of the type vector:

$$
\bar{U}_t(\theta_{1..t}) = \sum_{\tau=1}^{t} u_t(\theta_\tau;\theta_{1..\tau}) + U_t(\theta_{1..t}).
$$

In addition, the dynamic mechanism after the symmetrization will satisfy the *payment-frontloading* and *symmetry* properties:

Definition B.1 (Payment-frontloading). *A dynamic mechanism is payment-frontloading, if*

$$
u_t(\theta_{1..t}) = 0 \text{ for } t < T \quad \text{and} \quad u_T(\theta_{1..T}) \geq 0. \qquad \text{(PF)}
$$

The property is a stronger version of eP-IR.

Definition B.2 (Symmetry condition). *A dynamic mechanism satisfies the symmetry condition, if for every $t < s$:*

$$
\begin{aligned}
& \text{if } \bar{U}_t(\theta_{1..t}) = \bar{U}_t(\theta'_{1..t}) \text{ then:} \\
& x_s(\theta_{1..t}, \theta_{t+1..s}) = x_s(\theta'_{1..t}, \theta_{t+1..s}) \quad \text{and} \quad p_s(\theta_{1..t}, \theta_{t+1..s}) = p_s(\theta'_{1..t}, \theta_{t+1..s}).
\end{aligned} \qquad \text{(Symm)}
$$

Lemma B.3 (Payment frontloading)**.** *For any mechanism (x_t, p_t) satisfying DIC and eP-IR there is a mechanism also satisfying DIC and eP-IR with same allocation and ex-post revenue such that the agent is charged her full surplus in all periods except the last one.*

Lemma B.4 (Symmetrization)**.** *Any dynamic mechanism satisfying DIC and PF can be transformed into a mechanism $(x_t, p_t)_t$ with at least the same welfare and at least the same revenue as the original dynamic mechanism, satisfying: (i) DIC; (ii) PF; (iii) Symm.*

At first glance, our symmetrization lemma resembles the promised utility framework of Thomas and Worrall [TW90], which can be viewed as a symmetrization of the mechanism with respect to the continuation utilities U_t. Their result can be viewed as an application of the Principle of Optimality of the Theory of Dynamic Programming [Ber00], which describes the structure of an optimal solution that can be obtained by solving an infinite-size dynamic program. The symmetrization obtained in [TW90] is insufficient for our needs. Our solution is to transform the optimization program to a different space and apply the Principle of Optimality to the transformed program.

Next we prove the frontloading and symmetrization lemmas leading to the proof of Lemma 3.2:

Proof of Lemma B.3. Given a mechanism (x_t, p_t) satisfying DIC and eP-IR define mechanism $(x_t, \tilde{p}_t)$ such that $\tilde{p}_t(\theta_{1..t}) = v(\theta_t, x_t(\theta_{1..t}))$ for $t < T$ and

$$\tilde{p}_T(\theta_{1..T}) = \sum_{t=1}^{T} p_t(\theta_{1..t}) - \sum_{t=1}^{T-1} v(\theta_t, x_t(\theta_{1..t})).$$

The mechanism clearly has the revenue as the original, since for any $\theta_{1..T}$ we have $\sum_{t=1}^{T} p_t(\theta_{1..t}) = \sum_{t=1}^{T} \tilde{p}_t(\theta_{1..t})$. Since the ex-post allocation and ex-post revenue are the same in the two mechanisms for every $\theta_{1..T}$, the ex-post utility should also be the same. In particular, it should always be non-negative and therefore eP-IR holds. Since DIC can be formulated in terms of ex-post utilities it also holds after the transformation. □

One important property that we will use heavily is that since $u_t(\theta_{1..t}) = 0$ for all $t < T$, the continuation utility U_t and the partially realized utility $\bar{U}_t$ become the same things.

Proof of Lemma B.4. By Lemma B.3 we can assume (x_t, p_t) is a payment frontloading mechanism. Let's first define property Sym_t:

$$\begin{aligned}&\text{if } \bar{U}_t(\theta_{1..t}) = \bar{U}_t(\theta'_{1..t}) \text{ then } \forall s \geq t,\\ &x_s(\theta_{1..t}, \theta_{t+1..s}) = x_s(\theta'_{1..t}, \theta_{t+1..s}) \text{ and } p_s(\theta_{1..t}, \theta_{t+1..s}) = p_s(\theta'_{1..t}, \theta_{t+1..s})\end{aligned} \tag{Sym$_t$}$$

We will show that Sym_t works for all t by induction. Precisely: we show that if (x_t, p_t) is payment-frontloading satisfying Sym_t for $t \leq \tau - 1$ then we can transform it in a payment frontloading mechanism with at least the same revenue such that Sym_t holds for all $t \leq \tau$.

For the inductive step, partition the set of all possible type vectors $\theta_{1..\tau}$ into classes with the same partially-realized utility, i.e.,

$$S_\tau(x) = \{\theta_{1..\tau} | \bar{U}_\tau(\theta_{1..\tau}) = x\}$$

Now, for each x choose $\theta^*_{1..\tau}(x) \in S_\tau(x)$ maximizing the expected welfare of future periods

$$W_t(\theta_{1..t}) = \mathbb{E}\left[\sum_{t=\tau+1}^{T} v(\theta_t, x_t(\theta_{1..\tau}, \theta_{\tau+1..t}))\right]$$

Now, we define mechanism $(\tilde{x}, \tilde{p})$ such that $\tilde{x}_t = x_t$ and $\tilde{p}_t = p_t$ for $t \leq \tau$. For $t > \tau$ we have:

$$\tilde{x}_t(\theta_{1..t}) = x_t(\tilde{\theta}_{1..\tau}, \theta_{\tau+1..t}) \text{ where } \tilde{\theta}_{1..\tau} = \theta^*_{1..\tau}(\bar{U}_\tau(\theta_{1..\tau}))$$

$$\tilde{p}_t(\theta_{1..t}) = p_t(\tilde{\theta}_{1..\tau}, \theta_{\tau+1..t}) \text{ where } \tilde{\theta}_{1..\tau} = \theta^*_{1..\tau}(\bar{U}_\tau(\theta_{1..\tau}))$$

Now we argue that $(\tilde{x}_t, \tilde{p}_t)$ has the desired properties:

- it is still a payment frontloading mechanism, since allocation and payments from each type vector of length t are replaced by the allocation and payments of another type vector of length t, so the agent still has zero utility in all steps except the last one.

- it is still eP-IR. Let $\tilde{u}_t(\theta_{1..t})$ be the period utility of mechanism $(\tilde{x}_t, \tilde{p}_t)$. Since it is still payment-frontloading, $\tilde{u}_t(\theta_{1..t}) = 0$ for all $t < T$. So it is enough to argue that $\tilde{u}_T(\theta_{1..T}) \geq 0$. By the transformation, there is another type vector $\theta'_{1..T-1}$ such that:

$$\tilde{u}_T(\theta_{1..T}) = v(\theta_t, \tilde{x}_t(\theta_{1..t})) - \tilde{p}_t(\theta_{1..t}) = v(\theta_t, x_t(\theta'_{1..t-1}, \theta_t)) - p_t(\theta'_{1..t-1}, \theta_t) \geq 0$$

  since the original mechanism is also eP-IR and payment frontloading.

- it is still DIC. For $t > \tau$, the DIC condition follows directly from the fact that the original mechanism is DIC. For $t = \tau$ we use the fact that:

$$\theta_\tau = \arg\max_{\hat{\theta}_\tau} u_t(\theta_\tau; \hat{\theta}_{1..\tau}) + U_t(\hat{\theta}_{1..\tau}) = \arg\max_{\hat{\theta}_\tau} \tilde{u}_t(\theta_\tau; \hat{\theta}_{1..\tau}) + \tilde{U}_t(\hat{\theta}_{1..\tau}),$$

where $\tilde{U}_t$ is the continuation utility of the transformed mechanism. This expression holds since: (1) we didn't change the period utility of period τ; (2) we were careful to change the mechanism to preserve partially-realized utilities; (3) since the mechanism was a payment frontloading mechanism, the partially realized utilities coincide with continuation utilities, so we are also preserving continuation utilities. Finally, for $t < \tau$ we can use the same argument. Since the continuation utilities of period τ are preserved and the period utilities between period t and τ are preserved, the continuation utility of period t is also preserved.

- condition Sym_t holds for $t = \tau$. This condition holds by design.
- condition Sym_t holds for $t \leq \tau - 1$: the condition is clearly true for $s \leq \tau$. For $s < \tau \leq u$, consider two type vectors $\theta'_{1..s}$ and $\theta''_{1..s}$ with the same continuation utility in the original mechanism. Those must have the same continuation utility in the new mechanism as well, since we argue the continuation utilities are preserved for $t \leq \tau$. By the induction hypothesis the allocation and payments must be the same in the original mechanism for $(\theta'_{1..s}, \theta_{s+1..u})$ and $(\theta''_{1..s}, \theta_{s+1..u})$ for any type vector $\theta_{s+1..u}$. Therefore $U_\tau(\theta'_{1..s}, \theta_{s+1..\tau}) = U_\tau(\theta'_{1..s}, \theta_{s+1..\tau}) := x$ which means that both types are in the class, i.e., $(\theta'_{1..s}, \theta_{s+1..\tau}), (\theta''_{1..s}, \theta_{s+1..\tau}) \in S_\tau(x)$. Therefore:

$$\tilde{x}_u(\theta'_{1..s}, \theta_{s+1..u}) = x_u(\theta^*_{1..\tau}(x), \theta_{\tau+1..u}) = \tilde{x}_u(\theta''_{1..s}, \theta_{s+1..u})$$
$$\tilde{p}_u(\theta'_{1..s}, \theta_{s+1..u}) = p_u(\theta^*_{1..\tau}(x), \theta_{\tau+1..u}) = \tilde{p}_u(\theta''_{1..s}, \theta_{s+1..u})$$

- the expected welfare doesn't decrease, since we always replace a suffix of the mechanism with one with at least the same expected welfare:

$$\text{Sw} = \mathbb{E}\left[\sum_{t=1}^{\tau} v(\theta_t, x_t(\theta_{1..t})) + W_\tau(\theta_{1..\tau})\right] \leq \mathbb{E}\left[\sum_{t=1}^{\tau} v(\theta_t, x_t(\theta_{1..t})) + W_\tau(\theta^*_{1..\tau}(\bar{U}_\tau(\theta_{1..\tau})))\right]$$
$$= \mathbb{E}\left[\sum_{t=1}^{\tau} v(\theta_t, \tilde{x}_t(\theta_{1..t})) + \tilde{W}_\tau(\theta_{1..\tau})\right] = \tilde{\text{Sw}}$$

- the expected revenue doesn't decrease, since expected revenue is the difference of expected welfare and expected utility and we argue that welfare doesn't decrease and the expected utility is the same.

□

Proof of Lemma 3.2. A direct consequence of Lemma B.4 is that we can write $x_t = x_t(\theta_t, \bar{U}_{t-1})$

and $p_t = p_t(\theta_t, \bar{U}_{t-1})$. Also, $\bar{U}_t = \bar{U}_t(\theta_t, \bar{U}_{t-1})$ because by the payment frontloading property

$$\bar{U}_t = \mathbb{E}[\sum_{s=t+1}^{T} v(\theta_s, x_t(\theta_{t..s}, \bar{U}_{t-1})) - p_s(\theta_{t..s}, \bar{U}_{t-1})|\theta_t].$$

This allows us to define a bank account mechanism as follows. First we define the balance:

$$b_t^B(\theta_{1..t}) = \bar{U}_t(\theta_{1..t}) - \mu_t \quad \text{for} \quad t < T \quad \text{and} \quad b_T^B(\theta_{1..T}) = -\mu_T = 0,$$

where $\mu_t = \min_{\theta_{1..t}} \bar{U}_t(\theta_{1..t})$ for $t < T$ and $\mu_T = 0$. It will be useful to notice that by Jensen's inequality $\mu_0 \geq \mu_1 \geq \ldots \geq \mu_T = 0$. The allocation is the same as the original mechanism $x_t^B(\theta_{1..t}) = x_t(\theta_{1..t})$ and payments are computed as follows:

$$p_t^B(\theta_{1..t}) = p_t(\theta_{1..t}) - b_t^B(\theta_{1..t}) + b_{t-1}^B(\theta_{1..t-1}).$$

Since there is a one-to-one mapping between $\bar{U}_t$ and b_t^B, allocations, payments, and bank account updates can be computed from the previous state of the bank accounts, i.e., $x_t^B(\theta_{1..t}) = x_t^B(\theta_t, b_{t-1}) = x_t^B(\theta_t, b_{t-1}^B(\theta_{1..t-1}))$ and same for payments p_t^B. It will be useful to notice that we set payments and balance in such a way that:

$$v(\theta_t, x_t) - p_t + b_t^B = \bar{U}_t - \mu_t \tag{$\diamondsuit$}$$

This is true because for $t < T$, the per period utility $v(\theta_t, x_t) - p_t$ is zero since the mechanism is payment-frontloading; for $t = T$, $b_T^B = -\mu_T = 0$, and $v(\theta_T, x_T) - p_T = \bar{U}_T$, since by the payment-frontloading property the agent has non-zero utility only in the last period.

We will use this fact to check that the mechanism is a valid bank account mechanism. First note that by design $b_t^B(\theta_{1..t})$ is always non-negative and $b_0^B = 0$. Now we only need to check conditions IC, BI, and BU.

Condition IC follows from the definition of p_t^B and the fact that the original mechanism is DIC, since the maximization problem in IC becomes the same optimization in DIC with an additional constant term. For $t = T$ this is trivial since $b_T^B(\theta_{1..t}) = 0$. For $t < T$ we have:

$$\begin{aligned} u_t^B(\theta_t; \theta_{1..t-1}, \hat{\theta}_t) &= v(\theta_t, x_t^B(\theta_{1..t-1}, \hat{\theta}_t)) - p_t^B(\theta_{1..t-1}, \hat{\theta}_t) \\ &= v(\theta_t, x_t(\hat{\theta}_t, \bar{U}_t)) - p_t(\theta_{1..t-1}, \hat{\theta}_t) + \bar{U}_t(\theta_{1..t-1}, \hat{\theta}_t) - \big(\bar{U}_{t-1}(\theta_{1..t-1}) + \mu_t - \mu_{t-1}\big) \end{aligned}$$

since the term $\bar{U}_{t-1}(\theta_{1..t-1}) + \mu_t - \mu_{t-1}$ is a constant in $\hat{\theta}_t$ and $\bar{U}_t(\theta_{1..t-1}, \hat{\theta}_t) = U_t(\theta_{1..t-1}, \hat{\theta}_t)$ by the payment frontloading property. To check condition BI, we apply equation ($\diamondsuit$):

$$\mathbb{E}_{\theta_t}[v(\theta_t, x_t^B) - p_t^B] = \mathbb{E}_{\theta_t}[v(\theta_t, x_t) - p_t + b_t^B - b_{t-1}^B] = \mathbb{E}_{\theta_t}[\bar{U}_t - \mu_t - (\bar{U}_{t-1} - \mu_{t-1})] = \mu_{t-1} - \mu_t \geq 0$$

This establishes BI since the outcome is a constant that just depends on t but not on the value of b_t^B. Now, for condition BU we again apply equation (♢):

$$b_{t-1}^B + v(\theta_t, x_t) - p_t^B = b_{t-1}^B + v(\theta_t, x_t) - p_t + b_t^B - b_{t-1}^B = \bar{U}_t - \mu_t \geq b_t^B$$

where the last inequality holds with equality for all $t < T$. □

C Different Notions of IC and IR

C.1 Stronger IR notions

The main body of the paper focuses on satisfying DIC and eP-IR and the main design goals. In Lemma 3.1 we argue that bank account mechanisms satisfy even stronger notions. There are various variations over those notions that we can satisfy by slightly changing the mechanism. For example Lemma 3.1 implies that we satisfy the following notion of expected individual rationality continuation:

$$\mathbb{E}[\textstyle\sum_{\tau=t}^{T} u_\tau(\theta_\tau; \theta_{1..\tau}) | \theta_{1..t-1}] \geq 0$$

The reader might ask whether this is possible to satisfy the same notion ex-post with respect to the t-th type θ_t. In other words, can we satisfy the following notion?

$$\mathbb{E}[\textstyle\sum_{\tau=t}^{T} u_\tau(\theta_\tau; \theta_{1..\tau}) | \theta_{1..t}] \geq 0$$

Notice they only differ in the conditioning of the expectations. This can be achieved by any bank account mechanism by changing the payment rule to:

$$\hat{p}_t(\theta_t, b) = p_t(\theta_t, b) + b_t(\theta_t, b) - b, \text{ for } t < T$$

$$\text{and } \hat{p}_T(\theta_T, b) = p_T(\theta_T, b) - b.$$

The reader can verify that the all properties studied are preserved under this notion. In fact, condition BU implies that the previous transformation satisfies the even stronger notion of ex-post per-period individual rationality. I.e., under the $\hat{p}_t$ payment rules the mechanism satisfies for all realization of types and all periods:

$$u_t(\theta_t; \theta_{1..t}) \geq 0 \qquad \text{(pp-IR)}$$

This transformation almost preserves non-clairvoyance. If the original mechanism was non-clairvoyant the new mechanism is what we call *quasi-non-clairvoyant*. A quasi-non-clairvoyant mechanism is the one that needs to be told when the last period is at that period so that it can tailor its allocation and payment to the fact that we are in the last period. This is exactly what is required to implement the previous transformation.

We know there is a mechanism that is per-period individual rationality, dynamic incentive compatible and quasi-non-clairvoyant. Can we get the previous combination with actual non-clairvoyance instead of quasi-non-clairvoyance? The answer is unfortunately no.

Lemma C.1. *Any revenue that can be obtained by a non-clairvoyant mechanism satisfies DIC and pp-IR can also be obtained by running a static individually rational and incentive compatible auction in each period.*

Proof. The proof follows directly from Lemma D.1 in the following section, which says that a non-clairvoyant DIC mechanism must also satisfy per-period incentive compatibility. □

C.2 Stronger IC notions

Similarly we can ask the same question about incentive compatibility. Can we achieve even stronger notions of incentive compatibility? For example, can we achieve a version of DIC that holds for every realization of types in future periods instead of in expectation over future periods? We call a mechanism super dynamic incentive compatible:

$$\forall \hat{\theta}_{1..t-1}, \theta_{t+1..T},\ \theta_t = \arg\max_{\hat{\theta}_t} u_{t..T}(\theta_{t..T}; \hat{\theta}_{1..t-1}, \hat{\theta}_t, \theta_{t+1..T}), \tag{super-DIC}$$

where $u_{t..t'}(\theta_{t..t'}; \hat{\theta}_{1..t'}) = \sum_{s=t}^{t'} u_s(\theta_s; \hat{\theta}_{1..s})$. Unfortunately this notion is too strong as shown in Lemma C.2 which is restated here for convenience:

Lemma C.2. *For the auction setting where $v(\theta_t, x_t) = \theta_t \cdot x_t$, any revenue that can be obtained in a mechanism satisfying super-DIC and eP-IR can be obtained by running a static individually rational and incentive compatible mechanism in each period.*

Proof. Consider the single period mechanism with allocation defined by $\hat{x}(\hat{\theta}) = x_1(\hat{\theta})$. By the super-DIC property, for every $\theta_{2..T}$ the payment rule $\hat{p}(\hat{\theta}) = p_1(\hat{\theta}) - u_{2..T}(\theta_{2..T}; \hat{\theta}, \theta_{2..T})$ implements $\hat{x}$. Since the payment rule $\hat{p}$ is determined from $\hat{x}$ up to a constant, the term

$u_{2..T}(\theta_{2..T}; \theta_1, \theta_{2..T})$ must be decomposable in a term that depends only on θ_1 and a term depending on $\theta_{2..T}$. Say:

$$u_{2..T}(\theta_{2..T}; \theta_{2..T}) = \alpha(\theta_1) + \beta(\theta_{2..T})$$

Since $u_{1..T} = u_1(\theta_1; \theta_1) + \alpha(\theta_1) + \beta(\theta_{2..T})$ is non-negative for every type profile we can adjust α and β such that $u_1(\theta_1; \theta_1) + \alpha(\theta_1) \geq 0$ for every θ_1 and $\beta(\theta_{2..T}) \geq 0$ for every $\theta_{2..T}$. We can then define the following mechanism:

- allocate in the first period using $x_1(\theta_1)$ and charge $p_1(\theta_1) - \alpha(\theta_1)$
- allocate in all other periods using $\mathbb{E}_{\theta_1}[x_t(\theta_1, \theta_{2..t})]$ and charge $\mathbb{E}_{\theta_1}[x_t(\theta_1, \theta_{2..t})]$ adding an extra charge of $\mathbb{E}_{\theta_1}[\alpha(\theta_1)]$ in the last period.

We obtain a mechanism that is single period incentive compatible and individually rational for the first period, and a mechanism satisfying super-DIC and eP-IR for periods 2 to T. Notice that the revenue is still the same.

By induction we can find a mechanism that runs a static auction in each period and has the same revenue as the original mechanism. □

D Proof of the Non-Clairvoyance Gap Theorems

D.1 Characterization of Non-Clairvoyant Mechanisms

We start by proving the following strong property of non-clairvoyant mechanisms.

Lemma D.1. *If $x_t(F_{1..t}, \theta_{1..t})$, $p_t(F_{1..t}, \theta_{1..t})$ are a non-clairvoyant mechanism satisfying DIC, and $U_{t,T}(F_{1..T}, \theta_{1..t})$ for $t < T$ is the continuation utility of the corresponding clairvoyant mechanism:*

$$U_{t,T}(F_{1..T}, \theta_{1..t}) = \mathbb{E}_{\theta_{t+1..T} \sim F_{t+1..T}}[\textstyle\sum_{s=t+1}^{T} v(\theta_s, x_s(F_{1..s}, \theta_{1..s})) - p_s(F_{1..s}, \theta_{1..s})]$$

then $U_{t,T}(F_{1..T}, \theta_{1..t})$ doesn't depend on $\theta_{1..t}$, i.e., $U_{t,T}(F_{1..T}, \theta_{1..t}) = U_{t,T}(F_{1..T}, \theta'_{1..t})$.

Proof. Fix $F_{1..T}$ and $\theta_{1..t}$. First we show that $U_{t,T}(F_{1..T}, \theta_{1..t-1}, \hat{\theta}_t)$ doesn't depend on $\hat{\theta}_t$. Define the single period mechanism for a buyer with valuation $\hat{\theta}_t \sim F_t$ that allocates according to $\hat{x}_t(\hat{\theta}_t) = x_t(F_{1..t}, \theta_{1..t-1}, \hat{\theta}_t)$ and charges payments according to $\hat{p}(\hat{\theta}_t) = p_t(F_{1..t}, \theta_{1..t-1}, \hat{\theta}_t) -$

$U_{t,T}(F_{1..T}, \theta_{1..t-1}, \hat{\theta}_t)$. By the fact that the dynamic mechanism is DIC this mechanism must be incentive compatible, so the payment rule is uniquely defined by the allocation rule up to a constant. Now, define an alternative payment rule $p'_t(\hat{\theta}_t) = p_t(F_{1..t}, \theta_{1..t-1}, \hat{\theta}_t)$. The mechanism defined by $\hat{x}_t, p'_t$ must also be incentive compatible since the clairvoyant mechanism corresponding to the prior distribution sequence $F_{1..t}$ is also DIC. Since those are two single-period incentive compatible mechanisms with the same allocation rule, the payment rule must differ by a constant. Thus, the difference $U_{t,T}(F_{1..T}, \theta_{1..t-1}, \hat{\theta}_t)$ can't depend on $\hat{\theta}_t$.

Now we use induction to show that $U_{t,T}(F_{1..T}, \theta_{1..t-1}, \hat{\theta}_t)$ doesn't depend on θ_{t-1}. Since we know $U_{t,T}$ doesn't depend on θ_t we indicate it by writing $U_{t,T}(F_{1..T}, \theta_{1..t-1})$. By definition:

$$U_{t-1,T}(F_{1..T}, \theta_{1..t-2}) = U_{t-1,t}(F_{1..t}, \theta_{1..t-2}) + \mathbb{E}_{\theta_t \in F_t}[U_{t,T}(F_{1..T}, \theta_{1..t-1})]$$

Since the last term doesn't depend on θ_t we can remove the expectation:

$$U_{t,T}(F_{1..T}, \theta_{1..t-1}) = U_{t-1,T}(F_{1..T}, \theta_{1..t-2}) - U_{t-1,t}(F_{1..t}, \theta_{1..t-2}).$$

Hence, $U_{t,T}(F_{1..T}, \theta_{1..t-1})$ doesn't depend on θ_{t-1}. Repeating the same argument we can show $U_{t,T}$ depends only on the distributions $F_{1..T}$. □

Now, in order to prove Lemma 5.2 we first prove a symmetrization lemma in the style of Lemma B.4. There are some important differences: instead of the partially realized utility used in Lemma B.4 we will use the utility observed so far, which is a quantity we have access to in non-clairvoyant mechanisms since it only involves the past. The second major difference is that it won't involve payment frontloading, since we have no access to the future. The reason we can get away without those is the stronger property satisfied by non-clairvoyant mechanism described in Lemma D.1.

Lemma D.2 (Non-Clairvoyant Symmetrization)**.** *Given a non-clairvoyant dynamic mechanism $x_t(F_{1..t}, \theta_{1..t})$, $p_t(F_{1..t}, \theta_{1..t})$, there is a non-clairvoyant mechanism $\tilde{x}_t(F_{1..t}, \theta_{1..t}), \tilde{p}_t(F_{1..t}, \theta_{1..t})$ with the same revenue for each sequence of prior distributions, i.e., for each $F_{1..T}$:*

$$\mathbb{E}_{\theta_{1..t} \sim F_{1..t}}[\textstyle\sum_{t=1}^T p_t(F_{1..t}, \theta_{1..t})] = \mathbb{E}_{\theta_{1..t} \sim F_{1..t}}[\textstyle\sum_{t=1}^T \tilde{p}_t(F_{1..t}, \theta_{1..t})].$$

satisfying the following symmetry property: if $\sum_{s=1}^{t} \tilde{u}_s(F_{1..s}, \theta_{1..s}) = \sum_{s=1}^{t} \tilde{u}_s(F_{1..s}, \theta'_{1..s})$ *then:*

$$\tilde{x}_{t'}(F_{1..t}, F_{t+1,..,t'}, \theta_{1..t}, \theta_{t+1..t'}) = \tilde{x}_{t'}(F_{1..t}, F_{t+1,..,t'}, \theta'_{1..t}, \theta_{t+1..t'})$$
$$\tilde{p}_{t'}(F_{1..t}, F_{t+1,..,t'}, \theta_{1..t}, \theta_{t+1..t'}) = \tilde{p}_{t'}(F_{1..t}, F_{t+1,..,t'}, \theta'_{1..t}, \theta_{t+1..t'})$$

Proof. To prevent notations from being too verbose, define $u_{1..t}(F_{1..t}, \theta_{1..t}) = \sum_{s=1}^{t} u_s(F_{1..s}, \theta_{1..s})$. Assume the symmetric property holds for $t < \tau$. We will construct a mechanism for which the symmetric property holds for any $t \leq \tau$. Define $\tilde{x}_t$ and $\tilde{p}_t$ as follows. For $t \leq \tau$, let $\tilde{x}_t = x_t$ and $\tilde{p}_t = p_t$. For $t > \tau$ define:

$$\tilde{x}_t(F_{1..t}, \theta_{1..t}) = \mathbb{E}_{\theta'_{1..\tau} \sim F_{1..\tau}} \left[x_t(F_{1..t}, \theta'_{1..\tau}, \theta_{\tau+1..t}) \middle| \, u_{1..\tau}(F_{1..\tau}, \theta_{1..\tau}) = u_{1..\tau}(F_{1..\tau}, \theta'_{1..\tau}) \right]$$
$$\tilde{p}_t(F_{1..t}, \theta_{1..t}) = \mathbb{E}_{\theta'_{1..\tau} \sim F_{1..\tau}} \left[p_t(F_{1..t}, \theta'_{1..\tau}, \theta_{\tau+1..t}) \middle| \, u_{1..\tau}(F_{1..\tau}, \theta_{1..\tau}) = u_{1..\tau}(F_{1..\tau}, \theta'_{1..\tau}) \right]$$

In other words, we replace the allocation and payments in periods $t > \tau$ by the expected allocation and payments for types $\theta'_{1..\tau}, \theta_{\tau+1..t}$ such that the total utility accrued by the buyer in periods $1..\tau$ is the same as for $\theta_{1..\tau}$. Now we argue that this mechanism still has the desired properties:

- it is still non-clairvoyant: this is clear by construction since at period t the mechanism is only a function of $F_{1..t}$ and $\theta_{1..t}$. Notice that it is crucial that we symmetrize using a quantity that we can measure with information available at period t.
- it is still eP-IR. To check this property let $\tilde{u}_t$ be the utility under the new mechanism, then if E is the event that $u_{1..\tau}(F_{1..\tau}, \theta_{1..\tau}) = u_{1..\tau}(F_{1..\tau}, \theta'_{1..\tau})$, then:

$$\begin{aligned}\tilde{u}_{1..T}(F_{1..T}, \theta_{1..T}) &= u_{1..\tau}(F_{1..\tau}, \theta_{1..\tau}) + \mathbb{E}_{\theta'_{1..\tau}}[\textstyle\sum_{s=\tau+1}^{T} u_s(F_{1..s}, \theta'_{1..\tau}, \theta_{\tau+1..s}) | E] \\ &= \mathbb{E}_{\theta'_{1..\tau}}[u_{1..\tau}(F_{1..\tau}, \theta'_{1..\tau}) + \textstyle\sum_{s=\tau+1}^{T} u_s(F_{1..s}, \theta'_{1..\tau}, \theta_{\tau+1..s}) | E] \geq 0\end{aligned}$$

- it is still DIC. The DIC condition holds for $t > \tau$ since at that point the mechanism is simply a distribution of mechanisms satisfying the DIC condition. For $t \leq \tau$, we will use Lemma D.1 to argue that the expression in the maximization problem remains the same. In the following expression we omit $F_{1..t}$ for clarity of presentation:

$$\tilde{u}_t(\theta_{1..t}) + \tilde{U}_t(\theta_{1..t}) = u_t(\theta_{1..t}) + \mathbb{E}\left[\textstyle\sum_{s=t+1}^{\tau} u_s(\theta_{1..s})\right] + \mathbb{E}_{\theta'_{1..t}}\left[U_\tau(\theta'_{1..\tau}) | E(\theta_{1..\tau})\right]$$

where $E(\theta_{1..\tau})$ is the event determining the set of $\theta'_{1..\tau}$ that we will condition on. This event is a function of $\theta_{1..\tau}$. However, by Lemma D.1, U_τ is a constant so the expectation and the event we are conditioning on are irrelevant, therefore we have:

$$\tilde{u}_t(\theta_{1..t}) + \tilde{U}_t(\theta_{1..t}) = u_t(\theta_{1..t}) + \mathbb{E}\left[\sum_{s=t+1}^{\tau} u_s(\theta_{1..s})\right] + U_\tau(\theta'_{1..\tau}) = u_t(\theta_{1..t}) + U_t(\theta_{1..t})$$

- the symmetry condition holds for $t = \tau$ by design.
- the symmetry condition holds for $t < \tau$ using an argument analogous to the one used in Lemma D.1.
- the expected revenue is the same for the following reasons (again we omit $F_{1..T}$):

  $$\mathbb{E}_{\theta_{1..T}}[\sum_{t=1}^{T} \tilde{p}_t(\theta_{1..t})] = \mathbb{E}_{\theta_{1..\tau}}[\sum_{t=1}^{\tau} p_t(\theta_{1..\tau})] + \mathbb{E}_{\theta_{1..T}}\mathbb{E}_{\theta'_{1..\tau}}[\sum_{t=\tau+1}^{T} p_t(\theta'_{1..\tau}, \theta_{\tau+1..T}) | u_{1..\tau}(\theta_{1..\tau}) = u_{1..\tau}(\theta'_{1..\tau})]$$

  which equals to the original revenue since the distributions of $\theta_{1..\tau}$ and $\theta'_{1..\tau}$ are the same.

□

The symmetrization condition is the main ingredient to show that all non-clairvoyant mechanisms are bank account mechanisms. The reader is invited to contrast how much simpler this proof is than the proof of its clairvoyant counterpart. In some sense Lemma D.1 already provides us with most of the proof:

Proof of Lemma 5.2. Assume x_t, p_t satisfy the conditions in the Non-clairvoyant Symmetrization Lemma (Lemma D.1). Define the bank balance as $b_t(F_{1..t}, \theta_{1..t}) = \sum_{s=1}^{t} u_t(F_{1..t}, \theta_{1..t})$. From symmetrization it is clear that x_t, p_t can be written as a bank account mechanism. The BI condition follows directly from Lemma D.1. With the current definition of bank accounts, condition BU becomes trivial: the first inequality follows from eP-IR and the second one holds with equality. □

D.2 Lower bound for non-clairvoyant mechanisms

We will in this section prove Theorem 5.1. For this, let's initially define two distributions defined by their cdfs and parameterized by a constant $\mu > 0$ to be defined later:

$$F_1(\theta) = \left(1 - e^{-\mu^2}\right) \frac{\theta\mu}{\theta\mu+1} \text{ for } \theta \leq e^{\mu^2} \quad \text{and} \quad F_1(\theta) = 1 \text{ otherwise}$$

$$F_2(\theta) = \left[1 - \frac{\epsilon}{\theta}\right]^+$$

We will consider two scenarios: the first is that there is a single item with distribution F_1 and the second is that there are two items: the first with distribution F_1 and the second with distribution F_2. It is instructive to start by computing what is the optimal clairvoyant dynamic mechanism in each of the settings. By Lemma 3.2 we can restrict our attention to bank account mechanisms.

Scenario 1: One item with distribution F_1. Since there is only one period, the optimal mechanism is Myerson's auction. For the single-buyer case, it can be described as the posted price mechanism at ρ maximizing $\rho(1 - F_1(\rho))$, which is $\rho = e^{\mu^2}$, and the revenue is:

$$\textsc{Rev}^*(F_1) = \rho(1 - F_1(\rho)) = 1 + \tfrac{1}{\mu} + O(e^{-\mu^2})$$

Scenario 2: Two items with distributions F_1 and F_2. Since the optimal mechanism can be described as a bank account mechanism, assume x_t, p_t is the optimal bank account mechanism. By condition BU the state of the bank account in the end of period 1 is at most u_1 which is at most e^{μ^2}. The mechanism in the second period can be described as spending some amount which is at most the balance from the account and running an IC and IR mechanism. Since the distribution F_2 is such that $\rho(1 - F_2(\rho)) = \epsilon$ for all ρ (i.e., it is an equal-revenue distribution), the revenue obtained from the second period is at most $b_1 + \epsilon \le u_1 + \epsilon$. So the total revenue is at most the welfare of the first period plus ϵ. In other words, an upper bound to optimal revenue is $\mathbb{E}_{\theta_1 \sim F_1}[\theta_1] + \epsilon$.

Now we exhibit a mechanism that achieves that much revenue. In the first period, the item is given for free to the buyer and we add her value for the item into her bank account. In the second period, we first spend the entire balance of the bank account and then post a price $p(b_1)$ satisfying condition BI. No matter what price we post, the revenue will be $b_1 + \epsilon$. Therefore, the expected revenue of this mechanism is:

$$\textsc{Rev}^*(F_1, F_2) = \mathbb{E}_{\theta_1 \sim F_1}[\theta_1] + \epsilon = 1 + \mu + \epsilon + O(\mu e^{-\mu^2})$$

Comparison of the two scenarios: We note that depending on whether there will be a second item or not, we do two completely different things for the first item. If there is no second item, we allocate the second item with very low probability and charge a very high price if it is allocated. If there is a second item, we always allocate the first item and don't charge any amount for it. A non-clairvoyant mechanism needs to aim at balancing those two extremes: it needs to allocate the first item such that, if there is no second item, the revenue is good enough compared to the optimal single-item auction. But it also needs to make sure the bank balance after the first period is large enough to allow for more freedom in allocating the second item.

Non-Clairvoyant Mechanism Consider now a non-clairvoyant mechanism and let $x_1(F_1, \theta_1), p_1(F_1, \theta_1)$ be the auction for the first item with distribution F_1. This auction must be incentive compatible and individually rational, so it must be a distribution over posted price mechanisms, say, we use a random posted price $\rho \sim G$. Therefore:

$$\textsc{Rev}^M(F_1) = \mathbb{E}_{\rho \sim G}[\rho(1 - F_1(\rho))]$$

and since every non-clairvoyant mechanism can be written as a bank account mechanism (Lemma 5.2), we can use the same argument as in the scenario 2 above to argue that:

$$\textsc{Rev}^M(F_1, F_2) \leq \epsilon + \mathbb{E}_{\rho \sim G}\left[\mathbb{E}[\theta_1 \cdot \mathbf{1}_{\theta_1 \geq \rho}]\right]$$

Now, we are ready to prove the lower bound theorem:

Proof of Theorem 5.1. Assume that the non-clairvoyant mechanism is an α-approximation to the clairvoyant benchmark and consider the setup with F_1 and F_2 described in this section, then:

$$\frac{2}{\alpha} = 2\min\left(\frac{\textsc{Rev}^M(F_1)}{\textsc{Rev}^*(F_1)}, \frac{\textsc{Rev}^M(F_1,F_2)}{\textsc{Rev}^*(F_1,F_2)}\right) \leq \frac{\textsc{Rev}^M(F_1)}{\textsc{Rev}^*(F_1)} + \frac{\textsc{Rev}^M(F_1,F_2)}{\textsc{Rev}^*(F_1,F_2)} \leq \mathbb{E}_{\rho \sim G}\left[\beta(\rho)\right] \leq \max_\rho\left[\beta(\rho)\right]$$

$$\text{where} \qquad \beta(\rho) := \frac{\rho(1-F_1(\rho))}{\textsc{Rev}^*(F_1)} + \frac{\epsilon + \mathbb{E}[\theta_1 \cdot \mathbf{1}_{\theta_1 \geq \rho}]}{\textsc{Rev}^*(F_1,F_2)}$$

The remainder of the proof is Calculus-heavy[2] and involves explicitly substituting the values of those expressions and evaluating the maximum of $\beta(\rho)$. Taking the limit as $\mu \to \infty$ will provide us the desired bound.

Denote $r_1 = 1/\textsc{Rev}^*(F_1)$ and $r_{12} = 1/\textsc{Rev}^*(F_1, F_2)$, then

$$\beta(\rho) = r_1\rho(1 - F_1(\rho)) + r_{12}\left(\epsilon + \int_\rho^{e^{\mu^2}} \theta \mathrm{d}F_1(\theta)\right).$$

Taking derivative of β,

$$\begin{aligned}\beta'(\rho) &= r_1(1 - F_1(\rho) - \rho F_1'(\rho)) - r_{12}\rho F_1'(\rho) \\ &= r_1 - (1 - e^{\mu^2})r_1\tfrac{\rho\mu}{\rho\mu+1} - (r_1 + r_{12})(1 - e^{\mu^2})\tfrac{\rho\mu}{(\rho\mu+1)^2}\end{aligned}$$

[2] We will omit some less important calculation details, and Taylor expansion will be repeatedly used.

Denote $\zeta = 1 - e^{\mu^2}$ and let $\beta'(\rho) = 0$,

$$r_1(1-\zeta)(\rho\mu+1)^2 - r_{12}\zeta(\rho\mu+1) + (r_1+r_{12})\zeta = 0$$
$$\implies \rho\mu + 1 = \tfrac{r_{12}\zeta}{2r_1(1-\zeta)}\left(1 \pm \sqrt{1 - \tfrac{4r_1(1-\zeta)}{r_{12}\zeta}\left(1+\tfrac{r_1}{r_{12}}\right)}\right)$$

Since

$$\tfrac{r_1}{r_{12}} = \tfrac{\textsc{Rev}^*(F_1,F_2)}{\textsc{Rev}^*(F_1)} = \tfrac{1+\mu+\epsilon+O(\mu e^{-\mu^2})}{1+1/\mu+O(e^{-\mu^2})} = \mu + \epsilon + o(1),$$

$\frac{4r_1(1-\zeta)}{r_{12}\zeta}(1+\frac{r_1}{r_{12}}) \approx 4\mu^2 e^{-\mu^2} \ll 1$. Hence $\beta'(\rho) = 0$ has two roots. Because $\beta'(0) = r_1 > 0$, the local maximum of $\beta(\rho)$ is reached at the smaller root:

$$\rho^*\mu + 1 = \tfrac{r_{12}\zeta}{2r_1(1-\zeta)}\left(1 - \left(1 - \tfrac{4r_1(1-\zeta)}{r_{12}\zeta}\left(1+\tfrac{r_1}{r_{12}}\right) + o(e^{-\mu^2})\right)\right) = 1 + \mu + \epsilon + o(1)$$
$$\implies \rho^* = 1 + o(1).$$

Therefore the maximum value is reached at either ρ^* or e^{μ^2}:

$$\max_\rho \beta(\rho) = \max(\beta(\rho^*), \beta(e^{\mu^2})) = 1 + 1/\mu + o(1/\mu) \leq 1 + 2/\mu, \text{ for sufficiently large } \mu.$$

Hence the lower bound of α is obtained, which is 2 as $\mu \to \infty$:

$$\alpha \geq \tfrac{2}{\max_\rho \beta(\rho)} \geq \tfrac{2}{1+2/\mu}.$$

□

D.3 Lower bound for forgetful non-clairvoyant mechanisms

We build on the technique developed in the previous section to show that no forgetful non-clairvoyant bank balance mechanism can obtain a better than 1/3-approximation to the clairvoyant benchmark. We start by defining three distributions, which are parameterized by positive real numbers $\mu > 1$ and $\epsilon > 0$.

$$G_1(\theta) = \left[1 - \tfrac{1}{\theta}\right]^+;$$
$$G_2(\theta) = \begin{cases} \left(1 - e^{-\mu^2}\right)\frac{\theta\mu}{\theta\mu+1} & \text{for } \theta \leq e^{\mu^2}, \\ 1 - 1/\mu\theta & \text{for } e^{\mu^2} < \theta \leq e^{\mu e^{\mu}}, \\ 1 & \text{otherwise;} \end{cases}$$
$$G_3(\theta) = \begin{cases} 1 - \epsilon + \epsilon\left[1 - \frac{1}{\theta}\right]^+ & \text{for } \theta \leq e^{e^{\mu^2}/\epsilon}, \\ 1 & \text{otherwise.} \end{cases}$$

We also use $\emptyset$ to denote the point mass distribution at zero. We define three scenarios:

$$\textbf{Scenario A} : \emptyset, G_2, \emptyset \qquad \textbf{Scenario B} : G_1, G_2, \emptyset \qquad \textbf{Scenario C} : \emptyset, G_2, G_3,$$

where scenarios A and C are similar to the scenarios used in the proof of Theorem 5.1. G_2 and G_3 here are slightly modified versions of the F_1 and F_2 in the proof of Theorem 5.1, respectively.

Consider any single-buyer forgetful non-clairvoyant bank account mechanism $\mathcal{M}$ for $T = 3$ periods. Using the Spend notation introduced in Section 4, the allocation and pricing rule in each period t can be described as $x_t(\theta_t, b_{t-1}), p_t(\theta_t, b_{t-1})$ where $p_t(\theta_t, b_{t-1}) = p'_t(\theta_t, b_{t-1}) + s_t(b_{t-1})$ and, for each fixed b_{t-1}, the pair $x_t(\theta_t, b_{t-1}), p'_t(\theta_t, b_{t-1})$ is a single-period IC and IR mechanism. By Myerson's lemma [Mye81], a single-buyer IC and IR mechanism is a distribution over posted prices, i.e., there is a random variable $r_t(b_{t-1})$ such that $x_t(\theta_t, b_{t-1}) = \mathbb{E}[\mathbf{1}\{r_t(b_{t-1}) \leq \theta_t\}]$ and $p'_t(\theta_t, b_{t-1}) = \mathbb{E}[r_t(b_{t-1}) \cdot \mathbf{1}\{r_t(b_{t-1}) \leq \theta_t\}]$.

The main idea in the proof of Theorem 5.4 will be to argue that if $r_2(b_1)$ puts enough probability mass on high prices, then the revenue in scenario B is high, but the revenue in the worse of scenarios A and C is low. Conversely if the revenue in the worse of scenarios A and C is high, then the mechanism must put enough probability mass on low prices and, therefore, the revenue of scenario B is low. Since the mechanism is forgetful, the distribution of r_2 when $b_1 = 0$ is the same in all scenarios. Therefore, we can define

$$\kappa = \Pr[r_2(0) > e^{\mu^2}].$$

Here we won't specify the distribution of r_2 when $b_1 > 0$. Instead, we argue that there is an upper bound of the revenue for any fixed κ. We will prove it in a sequence of technical lemmas. To avoid repetition, it is implicit in the following statements and proofs that all mechanisms mentioned are forgetful bank account mechanisms.

Lemma D.3. *The ratio between* $\textsc{Rev}^{\mathcal{M}}(G_1, G_2, \emptyset)$ *and* $\textsc{Rev}^*(G_1, G_2, \emptyset)$ *is at most* κ *in the limit* $\mu \to \infty$.

Proof. We can decompose the revenue of any mechanism into three parts: $\mathbb{E}[p'_1(\theta_1) + s_2(b_1(\theta_1)) + p'_2(\theta_2, b_1(\theta_1))]$. The distributions G_1 and G_2 are chosen such that $\mathbb{E}[p'_1(\theta_1)] \leq 1$ and

$$\mathbb{E}[p'_2(\theta_2, b_1(\theta_1))] \leq 1 + 1/\mu + O(e^{-\mu^2}).$$

As for the spend part, by condition (BI), $s_2(b_1(\theta_1))$ is bounded by

$$s_2(b_1) \le \mathbb{E}\left[(\theta_2 - r_2(b_1))^+\right] - \mathbb{E}\left[(\theta_2 - r_2(0))^+\right] \le \mathbb{E}[\theta_2] - \mathbb{E}\left[(\theta_2 - r_2(0))^+\right] = \mathbb{E}[\min(\theta_2, r_2(0))].$$

In particular, with probability $1-\kappa$, $r_2(0) \le e^{\mu^2}$ and $\mathbb{E}[\min(\theta_2, r_2(0))|r_2(0)] \le \mathbb{E}[\min(\theta_2, e^{\mu^2})] = 1+\mu+O(\mu e^{-\mu^2})$, by the proof of Theorem 5.1. Otherwise, with probability κ, $e^{\mu^2} < r_2(0) \le \bar{\theta}_2 = e^{\mu e^{\mu}}$ and

$$\mathbb{E}[\min(\theta_2, r_2(0))|r_2(0)] \le \mathbb{E}[\theta_2] = 1 + \mu + O(\mu e^{-\mu^2}) + \frac{1}{\mu}\left(\ln \bar{\theta}_2 - \ln e^{\mu^2}\right) = 1 + e^{\mu} + O(\mu e^{-\mu^2}).$$

In meanwhile, by condition (BU), we have $s_2(b_1) \le b_1$ and $b_1(\theta_1) \le \theta_1$. Hence,

$$\mathbb{E}[s_2(b_1(\theta_1))] \le \mathbb{E}\left[\min\left(b_1(\theta_1), \mathbb{E}[\min(\theta_2, r_2(0))]\right)\right] \le \kappa\mu + (1-\kappa)\ln\mu + O\left(\frac{\ln\mu}{\mu}\right).$$

Thus,

$$\textsc{Rev}^{\mathcal{M}}(G_1, G_2, \emptyset) \le \kappa\mu + (1-\kappa)\ln\mu + 3 + O\left(\tfrac{\ln\mu}{\mu}\right) = \kappa\mu + O(\ln\mu).$$

On the other hand, in the optimal mechanism, the upper bound of the expected spend can be achieved. By setting $r_2^*(0)$ to the highest possible value $\bar{\theta}_2$, the surplus of the first period can be optimally extracted, i.e., let $r_2^*(b_1^*)$ be the value such that $\mathbb{E}\left[(\theta_2 - r_2^*(b_1^*))^+\right] = \min(b_1^*, \mathbb{E}[\theta_2])$. Then, according to condition (BI),

$$s_2^*(b_1^*) = s_2^*(0) + \mathbb{E}\left[(\theta_2 - r_2^*(b_1^*))^+\right] - \mathbb{E}\left[(\theta_2 - r_2^*(0))^+\right] = \min(b_1^*, \mathbb{E}[\theta_2]).$$

By letting $b_1^*(\theta_1) = (\theta_1 - 1)^+$, the expected spend can achieve

$$\mathbb{E}[s_2^*(b_1^*(\theta_1))] = \mathbb{E}[\min((\theta_1 - 1)^+, \mathbb{E}[\theta_2])] = \ln(1 + e^{\mu} + O(\mu e^{-\mu^2})) \ge \mu.$$

Hence, $\textsc{Rev}^*(G_1, G_2, \emptyset) \ge \mathbb{E}[s_2^*(b_1^*(\theta_1))] \ge \mu$ and

$$\frac{\textsc{Rev}^{\mathcal{M}}(G_1, G_2, \emptyset)}{\textsc{Rev}^*(G_1, G_2, \emptyset)} \le \frac{\kappa\mu + O(\ln\mu)}{\mu}$$

converging to κ as $\mu \longrightarrow \infty$. □

Lemma D.4. *If F_1 is the distribution defined in the proof of Theorem 5.1, then for any mechanism $\mathcal{M}$, $\left|\textsc{Rev}^{\mathcal{M}}(\emptyset, G_2, \emptyset) - \textsc{Rev}^{\mathcal{M}}(\emptyset, F_1, \emptyset)\right| \to 0$ as $\mu \to \infty$.*

Proof. Since only one period is non-trivial, the mechanism $\mathcal{M}$ is a single period IC and IR mechanism and therefore a distribution over the posted price mechanisms. So it is enough to show the lemma for the mechanism that posts a single price r.

Note that G_2 and F_1 are the same for $\theta \in [0, e^{\mu^2}] \cup (e^{\mu e^{\mu}}, +\infty)$. Therefore, for any posted prices no more than e^{μ^2} or larger than $e^{\mu e^{\mu}}$, the revenues in both cases are the same. For posted prices larger than e^{μ^2} but no more than $e^{\mu e^{\mu}}$, the revenue for $(\emptyset, G_2, \emptyset)$ is $r(1 - G_2(r)) = 1/\mu$, while the revenue for $(\emptyset, F_1, \emptyset)$ is 0. Their difference $1/\mu$ vanishes as μ goes to infinity. $\square$

Lemma D.5. *If F_1 is the distribution defined in the proof of Theorem 5.1, then for any mechanism $\mathcal{M}$, $\left|\text{Rev}^{\mathcal{M}}(\emptyset, G_2, G_3) - \text{Rev}^{\mathcal{M}}(\emptyset, F_1, G_3)\right| \to 0$ as $\mu \to \infty$.*

Proof. The revenue of each mechanism can be decomposed into $\mathbb{E}[p'_2(\theta_2, 0) + s_3(b_2(\theta_2)) + p'_3(\theta_3, b_2(\theta_2))]$. We compare each term separately: $|\mathbb{E}_{\theta_2 \sim G_2}[p'_2(\theta_2, 0)] - \mathbb{E}_{\theta_2 \sim F_1}[p'_2(\theta_2, 0)]| \leq 1/\mu \to 0$ by the previous lemma. For the second term, we know that the cdf of G_2 and F_1 differ only for $\theta_2 > e^{\mu^2}$, which happens with probability $1/(\mu e^{\mu^2})$ in either distribution. Since by (BI), the spend $s_3(b_2) \leq \mathbb{E}_{v_3 \sim G_3}[v_3] = \int_0^\infty (1 - G_3(v))\mathrm{d}v = \epsilon + e^{\mu^2}$, then

$$|\mathbb{E}_{\theta_2 \sim G_2}[s_3(b_2)] - \mathbb{E}_{\theta_2 \sim F_3}[s_3(b_2)]| \leq \frac{\epsilon + e^{\mu^2}}{\mu e^{\mu^2}} \to 0 \text{ as } \mu \to 0.$$

We use the same idea for the last term: since $0 \leq \mathbb{E}_{\theta_3 \sim G_3}[p'_3(\theta_3, b_2(\theta_2))] \leq \epsilon$ and the cdf of G_2 and F_1 differ only for $\theta_2 > e^{\mu^2}$,

$$|\mathbb{E}_{\theta_2 \sim G_2, \theta_3 \sim G_3}[p'_3(\theta_3, b_2(\theta_2))] - \mathbb{E}_{\theta_2 \sim F_3, \theta_3 \sim G_3}[p'_3(\theta_3, b_2(\theta_2))]| \leq \frac{\epsilon}{\mu e^{\mu^2}} \to 0 \text{ as } \mu \to 0.$$

$\square$

Lemma D.6. *Let F_1 and F_2 be the distributions defined in the proof of Theorem 5.1. For any mechanism $\mathcal{M}$, there is another mechanism $\tilde{\mathcal{M}}$ such that*

$$\text{Rev}^{\mathcal{M}}(\emptyset, F_1, G_3) \leq \text{Rev}^{\tilde{\mathcal{M}}}(\emptyset, F_1, F_2) \leq \text{Rev}^{\mathcal{M}}(\emptyset, F_1, G_3) + \epsilon.$$

Proof. We prove the case where the mechanism (x_3, p'_3) is a posted price mechanism with price $r_3(b_2)$. Since (x_3, p'_3) is an IC and IR single-buyer mechanism, it is a distribution over posted prices and hence, the same proof holds by treating $r_3(b_2)$ as a random price instead of a single price.

Since the balance of $\mathcal{M}$ at the end of the second period b_2 can be at most e^{μ^2}, then according to (BI) and (BU), we have

$$\forall b_2,\ r_3(0) \geq r_3(b_2) \quad \text{and} \quad \mathbb{E}_{v_3 \sim F_2}[(v_3 - r_3(b_2))^+] - \mathbb{E}_{v_3 \sim F_2}[(v_3 - r_3(0))^+] \leq b_2 \leq e^{\mu^2}. \tag{D.1}$$

Hence, there always exists $\tilde{r}_3(b_2) \geq 0$, such that

$$\forall b_2, \mathbb{E}_{v_3 \sim G_3}[(v_3 - \tilde{r}_3(b_2))^+] = \mathbb{E}_{v_3 \sim F_2}[(v_3 - r_3(b_2))^+] - \mathbb{E}_{v_3 \sim F_2}[(v_3 - r_3(0))^+],$$

because the right-hand-side is at most e^{μ^2} and the left-hand-side could be as large as e^{μ^2} when $\tilde{r}_3 = 1$. Therefore, $\mathbb{E}_{v_3 \sim G_3}[(v_3 - 1)^+] = \epsilon(\ln e^{e^{\mu^2}/\epsilon} - \ln 1) = e^{\mu^2}$.

Therefore, we can let $\tilde{\mathcal{M}}$ be the same as $\mathcal{M}$ except when $\tilde{r}_3(0) = e^{e^{\mu^2}/\epsilon}$ and $\tilde{r}_3(b_2)$ is defined according to (D.1). Then the spend terms in both mechanisms are always the same, i.e., $\tilde{s}_3(b_2) = s_3(b_2)$, and the expected direct payments are: $0 \leq \mathbb{E}[p_3] \leq \epsilon$ and $\mathbb{E}[p'_3] = \epsilon$. $\square$

Lemma D.7. *In the limit when $\mu \to 0$ and $\epsilon \to 0$, the minimum*

$$\min\left(\frac{\text{Rev}^{\mathcal{M}}(\emptyset, G_2, \emptyset)}{\text{Rev}^*(\emptyset, G_2, \emptyset)}, \frac{\text{Rev}^{\mathcal{M}}(\emptyset, G_2, G_3)}{\text{Rev}^*(\emptyset, G_2, G_3)}\right) \leq \frac{1-\kappa}{2}.$$

Proof. If $r_2(0) > e^{\mu^2}$, then $\text{Rev}^{\mathcal{M}}(\emptyset, G_2, \emptyset) \leq 1/\mu \to 0$ while $\text{Rev}^*(\emptyset, G_2, \emptyset) \to 1$ once $\mu \to \infty$. In the remaining case, $r_2(0) \leq e^{\mu^2}$, we use the previous lemmas to argue that:

$$\min\left(\frac{\text{Rev}^{\mathcal{M}}(\emptyset, G_2, \emptyset)}{\text{Rev}^*(\emptyset, G_2, \emptyset)}, \frac{\text{Rev}^{\mathcal{M}}(\emptyset, G_2, G_3)}{\text{Rev}^*(\emptyset, G_2, G_3)}\right) \leq \min\left(\frac{\text{Rev}^{\mathcal{M}}(\emptyset, F_1, \emptyset)}{\text{Rev}^*(\emptyset, F_1, \emptyset)}, \frac{\text{Rev}^{\mathcal{M}}(\emptyset, F_1, F_2) + \epsilon}{\text{Rev}^*(\emptyset, F_1, F_2) - \epsilon}\right).$$

By the proof of Theorem 5.1, $\text{Rev}^*(\emptyset, F_1, F_2) - \epsilon = 1 + \mu + O(\mu e^{-\mu^2}) > 0$ and hence, the right-hand-side above can be bounded by $1/2 + 1/\mu$ for $\epsilon \to 0$, where the bound converges to $1/2$ as $\mu \to \infty$.

Combining the two observations, we get the statement of the lemma. $\square$

Proof of Theorem 5.4. Combining Lemma D.3 and Lemma D.7, we observe that for any mechanism $\mathcal{M}$,

$$\max_{\mathcal{M}} \min_{G_{1..T}} \text{Rev}^{\mathcal{M}}(G_{1..T})/\text{Rev}^*(G_{1..T}) \leq \max_{0 \leq \kappa \leq 1} \min\left(\kappa, \frac{1-\kappa}{2}\right) = \frac{1}{3}.$$

The NonClairvoyantBalance mechanism achieves that bound (by Theorem 4.2) and therefore, is maximin optimal. $\square$

E Details Omitted from Multiple Buyers

E.1 Multi-buyer dynamic mechanism design

We start by extending the concepts in the paper to multiple buyers. Consider a set N of n agents who participate in the mechanism for T periods. For each agent $i \in N$ and each $t \in \{1, \ldots, T\}$ the type θ_t^i of agent i in period t is drawn independently from a distribution F_t^i. When we omit the superscript i we refer to the vector of types $\theta_t = (\theta_t^1, \ldots, \theta_t^n)$. As usual in mechanism design we refer to θ_t^{-i} as the vector of types of all agents except i. Agent i has a value $v^i : \Theta \times \mathcal{O} \to \mathbb{R}_+$. A dynamic mechanism corresponds to pairs of maps:

- Outcome: $x_t : \Theta^{tN} \times (\Delta\Theta)^{TN} \to \mathcal{O}$
- Payment: $p_t : \Theta^{tN} \times (\Delta\Theta)^{TN} \to \mathbb{R}$

Similarly to the single buyer case, we can define the notion of continuation utility $U_t^i(\hat{\theta}_{1..t}; F_{1..T})$ as the expected total utility of a buyer in periods $t+1$ to T if her history of reports up to period t is $\hat{\theta}_{1..t}$ and all the buyers report truthfully from period $t+1$ onwards. This allows us to define the analogue of condition DIC for multiple buyers, which we call Dynamic Bayesian Incentive Compatibility. We call it Bayesian since each buyer takes expectations over the behavior of all other buyers assuming they bid truthfully. The condition can be written as follows:

$$\theta_t = \arg\max_{\hat{\theta}_t} \mathbb{E}_{\theta_t^{-i}} \left[u_t^i(\theta_t^i; \hat{\theta}_{1..t-1}, (\theta_t^{-i}, \hat{\theta}_t^i)) + U_t^i(\hat{\theta}_{1..t-1}, (\theta_t^{-i}, \hat{\theta}_t^i)) \right] \tag{DBIC}$$

We recall that while the condition DIC for a single buyer can be justified by the dynamic version of the revelation principle, no such equivalence can be obtained for multiple buyers. What we have here is an ex-post incentive compatibility: it is optimal for a buyer to report her type truthfully as long as all the other buyers also do so. We refer to [AS13] or [PST14] for a discussion of the relation between incentive compatibility in dynamic settings and the revelation principle, as well as [MR92] and [BM05] for the comparison of dominant-strategy implementation, ex-post implementation, and Bayesian implementation.

The condition eP-IR is generalized in the natural way. Every buyer derives non-negative utility in every sample path if she is behaving truthfully.

The notion of non-clairvoyance corresponds again to the same measure theoretic restriction that the allocation and payment functions in time t must be measurable with respect to $(\theta_{1..t}, F_{1..t})$, i.e., can't depend on distributional knowledge of future periods.

E.2 Multi-buyer bank account mechanisms

We define a bank account mechanism for n buyers as:

- A static single-period mechanism $x_t^B(\theta_t, b), p_t^B(\theta_t, b)$ parameterized by an n-dimensional bank balance $b \in \mathbb{R}_+^n$ that is single-period Bayesian incentive compatible, i.e., satisfies the multi-buyer version of IC and satisfies the multi-buyer version of BI:

$$\mathbb{E}_{\theta_t}[v^i(\theta_t^i; x_t^B(\theta_t, b)) - p_t^B(\theta_t, b)] \text{ is a non-negative constant not depending on } b.$$

- A balance update policy $b_t^B(\theta_t, b)$ satisfying a multi-buyer equivalent of condition BU:

$$0 \leq b_t^{B,i}(\theta_t, b) \leq b^i + u_t^{B,i}(\theta_t, b) \text{ and } b_0^i = 0$$

As before, it is useful to define a notion of spend s_t^i as follows:

$$s_t^i(b_{t-1}) = \left[-\min_{\theta_t^i} \mathbb{E}_{\theta_t^{-i}}\left[v^i(\theta_t^i, x_t(\theta_t, b_{t-1})) - p_t^i(\theta_t, b_{t-1})\right]\right]^+$$

Both the clairvoyant (Lemma 3.2) and non-clairvoyant (Lemma 5.2) reductions still hold in the multi-buyer setting with essentially the same proofs by adapting the notation.

E.3 Proof of Theorem 6.1

Besides the notion of the spend, it will also be useful to define an auxiliary notion called the deposit,

$$d_t^i(\theta_t, b_{t-1}) = b_t^i(\theta_t, b_{t-1}) - b_{t-1}^i + s_t^i(b_{t-1}),$$

so that we can describe the balance update policy in terms of the spend and deposit:

$$b_t^i = b_{t-1}^i + d_t^i - s_t^i.$$

In particular, if we write the p_t^i as $p_t^i = {p'}_t^i + s_t^i$ and $u_t^i = {u'}_t^i - s_t^i$, then the (BU) condition can be rewritten as

$$d_t^i \leq {u'}_t^i.$$

Proof of Theorem 6.1. Fix a time horizon T and distributions F_t^i for $t = 1..T$ and $i = 1..n$. Let (x^*, p^*) be the optimal clairvoyant mechanism for this setting. By the multi-buyer version of Lemma 3.2, we can write the bank account mechanism in terms of a spend policy s_t^*, a deposit policy d_t^*, and an IC and IR payment function ${p'}_t^*$ such that:

$$p_t^{*i} = {p'}_t^{*i} + s_t^{*i} \qquad b_t^{*i} = b_{t-1}^{*i} - s_t^{*i} + d_t^{*i}.$$

Similarly, let x_t, p'_t, s_t, d_t describe the NonClairvoyantBalance mechanism where the spend term corresponds to the expected utility of the Money Burning.

Step 1: Bounding p'^* using the Myerson component. Our first observation is that since for each period x^*_t, p'^*_t is individually rational and Bayesian incentive compatible, its revenue must be dominated by the Myerson auction: $\mathbb{E}_{\theta_t}[\sum_i p'^{*i}_t(\theta_{1..t})] \leq \mathbb{E}_{\theta_t}[\sum_i p^{M,i}_t(\theta_t)]$. This already tells us that the revenue we obtain from selling 1/5 fraction of each item using Myerson's auction dominates within a factor of 5 the $\mathbb{E}[\sum_{i,t} p'^{*i}_t]$ component of the revenue of the optimal clairvoyant mechanism.

Step 2: Lower bound to the balance of the non-clairvoyant mechanism. We are left to show that the remaining component $\mathbb{E}[\sum_{i,t} s^{*i}_t]$ of the revenue of the optimal clairvoyant mechanism is dominated by the combination of the Second Price Auction and the Money Burning Auction within a factor of 5. We will show by induction that for every fixed sequence of types and for all buyers $\theta_{1..T}$ the following invariant holds. Since the types for all buyers are fixed for all periods, we will omit the type vectors in the notation.

$$b^i_t + \sum_{\tau=1}^t s^i_\tau \geq \tfrac{2}{5}(b^{*i}_t + \sum_{\tau=1}^t s^{*i}_\tau - \sum_{\tau=1}^t \theta^{(2)}_\tau x^{*i}_\tau) \tag{E.1}$$

where $\theta^{(2)}_\tau$ is the second highest type. This is true for $t = 0$ since both balances are initially zero. Now, assume it is valid for t then substituting the balance update formula $b^i_{t+1} = b^i_t - s^i_{t+1} + d^i_{t+1}$ for both the non-clairvoyant and the clairvoyant mechanism we obtain:

$$b^i_{t+1} + \sum_{\tau=1}^{t+1} s^i_\tau - d^i_{t+1} \geq \tfrac{2}{5}(b^{*i}_{t+1} + \sum_{\tau=1}^{t+1} s^{*i}_\tau - \sum_{\tau=1}^t \theta^{(2)}_\tau x^{*i}_\tau - d^{*i}_{t+1}).$$

By (BU), $d^{*i}_{t+1} \leq u'^{*i}_{t+1} \leq \theta^i_{t+1} x^{*i}_{t+1}$. If i is not the agent with the highest type then $\theta^i_{t+1} \leq \theta^{(2)}_{t+1}$ and we are done by the fact that $d^i_{t+1} \geq 0$ and $\theta^{(2)}_{t+1} x^{*i}_{t+1} \geq d^{*i}_{t+1}$. If i is the agent with the highest type, then

$$d^i_{t+1} = \tfrac{2}{5}(\theta^i_{t+1} - \theta^{(2)}_{t+1}) \geq \tfrac{2}{5}(\theta^i_{t+1} - \theta^{(2)}_{t+1}) x^{*i}_{t+1} \geq \tfrac{2}{5}(d^{*i}_{t+1} - \theta^{(2)}_{t+1} x^{*i}_{t+1}),$$

since we only deposit in the Second Price Auction mechanism for the top agent. Substituting this bound we obtain the invariant for $t + 1$.

Step 3: Charging scheme for spend. We will construct a charging scheme to re-attribute the spends of the non-clairvoyant mechanism in a way that makes it resemble more the spends of the optimal clairvoyant mechanism. For each fixed $\theta_{1..T}$ we will define a charging scheme $c_t^i \geq 0$ such that for each period t we have $\sum_i c_t^i \leq \sum_i s_t^i$. We will do so in such a way that we can more easily compare s_t^{*i} with c_t^i.

We know by (BI) that there is a solution to the Money Burning problem in period t with $\mathbb{E}[\tilde{u}_t^i] \geq s_t^{*i}$ since the clairvoyant mechanism with balance b_{t-1}^* provides such a solution. Thus, by rescaling the mechanism there must be a solution to the money burning problem with constraints $\mathbb{E}[\tilde{u}_t^i] \leq \frac{5}{2} b_{t-1}^i$ such that $\mathbb{E}[\tilde{u}_t^i] = \min\left(s_t^{*i}, \frac{5}{2} b_{t-1}^i\right)$. In particular it means:

$$\sum_i s_t^i \geq \frac{2}{5} \sum_i \min\left(s_t^{*i}, \frac{5}{2} b_{t-1}^i\right)$$

This motivates to define the following charging scheme:

$$c_t^i = \min\left(\frac{2}{5} s_t^{*i}, b_{t-1}^i\right).$$

Based on how we compute the charge we divide the set of agents in each period in a set A_t of agents ahead and a set B_t of agents behind. We say agent i is behind ($i \in B_t$) if $b_{t-1}^i \leq \frac{2}{5} s_t^{*i}$ and we say that i is ahead ($i \in A_t$) otherwise. For $i \in B_t$ we can produce a good bound on the total spend using (E.1):

$$c_t^i = b_{t-1}^i \geq \frac{2}{5}(b_{t-1}^{*i} + \sum_{\tau=1}^{t-1} s_\tau^{*i} - \sum_{\tau=1}^{t-1} \theta_\tau^{(2)} x_\tau^{*i}) - \sum_{\tau=1}^{t-1} s_\tau^i.$$

Re-organizing the expression and using that $s_t^{*i} \leq b_{t-1}^{*i}$ we get:

$$c_t^i + \sum_{\tau=1}^{t-1} s_\tau^i + \frac{2}{5} \sum_{\tau=1}^{t-1} \theta_\tau^{(2)} x_\tau^{*i} \geq \frac{2}{5} \sum_{\tau=1}^{t} s_\tau^{*i} \tag{E.2}$$

A similar bound can be used to bound an ahead agent $i \in A_t$. Let t' be the last period before t where $i \in B_{t'}$. This is well-defined since all agents are behind in period zero. Therefore (E.2) holds for t'. Now, we can sum $\sum_{\tau=t'+1}^{t} c_\tau^i \geq \frac{2}{5} \sum_{\tau=t'+1}^{t} s_\tau^{*i}$ to that bound and get:

$$\sum_{\tau=1}^{t'-1} s_\tau^i + \sum_{\tau=t'}^{t} c_\tau^i + \frac{2}{5} \sum_{\tau=1}^{t'-1} \theta_\tau^{(2)} x_\tau^{*i} \geq \frac{2}{5} \sum_{\tau=1}^{t} s_\tau^{*i} \tag{E.3}$$

Step 4: Bounding the spend of the non-clairvoyant mechanism. Either if $i \in B_t$ (E.2) or $i \in A_t$ (E.3) we can bound the spend as follows:

$$\sum_{\tau=1}^{t} s_\tau^i + \sum_{\tau=1}^{t} c_\tau^i + \frac{2}{5} \sum_{\tau=1}^{t} \theta_\tau^{(2)} x_\tau^{*i} \geq \frac{2}{5} \sum_{\tau=1}^{t} s_\tau^{*i}.$$

Summing over all agents i and using the fact that $\sum_i c_t^i \leq \sum_i s_t^i$ we have:

$$2 \sum_i \sum_{\tau=1}^{T} s_\tau^i + \frac{2}{5} \sum_{\tau=1}^{t} \theta_\tau^{(2)} \geq \frac{2}{5} \sum_i \sum_{\tau=1}^{t} s_\tau^{*i}.$$

Dividing the expression by 2, we see that the sum of total spends of the non-clairvoyant mechanism together with the revenue obtained from the second price auction component gives us a 5-approximation to the total spend of the optimal clairvoyant mechanism. □

E.4 Other omitted proofs from Section 6

Proof of Lemma 6.2. By the BI property, the expected utility in subsequent rounds is not a function of the current reported type, so it is enough to argue that the three components of the NonClairvoyantBalance mechanism are dominant strategy incentive compatible in the static sense. This is trivial to check for the second price and Myerson components. For the Money Burning auction, we refer the reader to Appendix F where we discuss how to construct this component. □

Proof of Theorem 6.3. The proof is almost implied by the arguments we made in the proof of Theorem 6.1. Since there are only 2 periods in total, then

- The spend of the clairvoyant mechanism in the first period is zero: $\sum_i s_1^{*i} = 0$. Therefore the non-clairvoyant mechanism doesn't lose any spend for not including Money Burning Auction in the first period.
- The total spend only depends on the balance from the first period (b_1). Therefore the non-clairvoyant mechanism doesn't lose any spend for not including Second Price Auction in the second period.
- The total spend only comes from the second period. Hence the Money Burning Auction is optimal in the spend.

Putting these observations together, we can conclude that for any type vector sequence $\theta_{1..2}$,

$$\theta_1^{(2)} + \sum_i s_1^i + s_2^i \geq \frac{1}{2} \sum_i s_1^{*i} + s_2^{*i}$$

Combining this with the fact that the non-clairvoyant mechanism sells half of the item via Myerson's Auction, we conclude that it is a non-clairvoyant 2-approximation. □

F Implementation of NonClairvoyantBalance

Here we show that all the three components of the NonClairvoyantBalance mechanism are simple auctions: each of them corresponds to maximizing some notion of virtual values.

The first component of the NonClairvoyantBalance mechanism is a Second Price Auction which doesn't use any information about the distribution and the virtual value is simply the buyer's value. The second component is the Myerson auction, which, of course, is a virtual value maximizer.

Most of our work will be focused on arguing that the third component — the Money Burning auction with utility constraints has a simple format and can be implemented as a virtual value maximizer. In what follows, for the ease of presentation, we will focus on discrete distributions. Assume therefore that the space of valuation functions is a finite set of non-negative numbers, i.e., $\Theta = \{\theta_1, \ldots, \theta_K\} \subset \mathbb{R}_+$. As we will focus on a single period, we ignore subscripts t. Instead θ_j will refer to the j-th value in support of the distribution. As before let n be the number of buyers. The distributions F^i will be discrete distributions represented by a vector of K non-negative numbers $f^i(\theta_1), \ldots f^i(\theta_K)$ summing to 1. We will also denote the cdf of the distribution by $F^i(\theta) = \sum_{\theta_j \leq \theta} f^i(\theta_j)$.

F.1 Optimal Money Burning with Caps is a scaled virtual value maximizer

Since the optimal Money Burning mechanism can be written as an optimization problem in the reduced form, it is possible to directly obtain an algorithm using the framework of Cai, Daskalakis, and Weinberg [CDW12a, CDW12b]. For the special case of money burning, an alternative solution goes through the techniques developed by Hartline and Roughgarden [HR08]. A black-box application of [CDW12a, CDW12b] guarantees that the auction is Bayesian Incentive Compatible. For Lemma 6.2 it will be useful to describe the auction via the virtual value technique of [HR08] to show that the optimal capped money burning auction is dominant strategy incentive compatible. We discuss the construction below.

Without any caps on the utilities the optimal money burning auction was analyzed by Hartline and Roughgarden [HR08] and shown to be a virtual value maximization for a different notion of virtual values known as *virtual values for utility*. As in the Myerson auction, the virtual values of [HR08] can be computed as a function of the distribution, and if not monotone they require to be ironed using the same procedure used to iron the

Myersonian virtual values. While originally developed for continuous distributions, the exact approach[3] described by Elkind [Elk07] can be used to compute ironed virtual values for utility for all buyers. We can summarize their result as follows:

Theorem F.1 (Hartline and Roughgarden)**.** *Given distributions $F^1, \ldots, F^n$ of support $\Theta = \{\theta_1, \ldots, \theta_K\}$, there exist non-decreasing maps $\vartheta^i : \Theta \to \mathbb{R}$ (called ironed virtual values for utility) such that for any Bayesian incentive compatible and individually rational mechanism (x^i, p^i) and for every agent i:*

$$\mathbb{E}_{\theta \sim F}[u^i(\theta)] = \mathbb{E}_{\theta \sim F}[\vartheta^i(\theta^i) x^i(\theta)]$$

Moreover, the optimal mechanism (with or without utility caps) is such that the allocation and payments only depend on the virtual values $\vartheta^i(\theta^i)$.

The proof of the theorem follows from combining Lemma 2.6, Lemma 2.8, and Theorem 2.9 in [HR08]. For the moreover part, even though their paper doesn't consider any sort of utility caps, the presence of caps doesn't affect any of their proofs.

From Theorem F.1 we can describe the optimal auction as monotone allocation that depends only on virtual values. We abuse notations and use f^i to denote the distribution on the virtual values, i.e., $f^i(\bar{\vartheta}^i) = \sum_{\theta^i \in \Theta; \vartheta^i(\theta^i) = \bar{\vartheta}^i} f^i(\theta^i)$. We also define the allocation directly in terms of virtual values $x^i(\vartheta)$. Now we are ready to describe the format of the optimal auction:

Theorem F.2 (Optimal Capped Money Burning)**.** *The auction maximizing capped utility $\sum_i \min(b^i, \mathbb{E}[u^i])$ is parameterized by w^i, q^i, which chooses the agent with largest scaled virtual value $w^i \vartheta^i$ (subject to some tie breaking rule) and allocates to this agent with probability q^i.*

Proof. Using Theorem F.1, we can formulate the optimal money burning with caps problem as finding a monotone allocation function $x^i(\vartheta)$ defined on the virtual values maximizing

[3] Given a discrete distribution described by $f(\theta_1), \ldots, f(\theta_K)$ non-negative and summing to 1 with $\theta_1 < \theta_2 < \ldots < \theta_K$, Elkind [Elk07] defines a discrete notion of the Myersonian virtual value as $\varphi^i_j = \theta_j - (\theta^i_{j+1} - \theta^i_j) \frac{1-F^i(\theta_j)}{f^i(\theta_j)}$. Those are then ironed by defining for each i a set of K 2-dimensional points $(F(\theta^i_j), \sum_{j' \leq j} f^i_{j'} \varphi^i_{j'})$, computing the lower convex hull and defining the ironed virtual values as the slopes of segments of the convex hull corresponding to each point. The same exact computation can be done by replacing the original Myersonian notion of virtual values φ^i_j with the definition of virtual values for utility $\vartheta^i_j = (\theta^i_{j+1} - \theta^i_j) \frac{1-F^i(\theta_j)}{f^i(\theta_j)}$.

$\mathbb{E}[\vartheta^i x^i(\vartheta)]$. We will solve the problem:

$$\max \textstyle\sum_i \min\left(b^i, \mathbb{E}[\vartheta^i x^i(\vartheta)]\right) \quad \text{s.t. monotonicity}$$

and will rescaled the x^i by multiplying it by a probability q^i such that it obeys the constraints $\mathbb{E}[\sum_i \vartheta^i x^i(\vartheta)] \le b^i$ while keeping the same objective value. In the following formulation we relax the constraint that the allocation needs to be monotone and obtain the following primal-dual pair:

$$\begin{aligned} \max_{x,u} \quad & \textstyle\sum_i u^i \\ \text{s.t.} \quad & u^i \le \textstyle\sum_\vartheta \vartheta^i x^i(\vartheta) f(\vartheta),\ \forall i \quad & (w^i) \\ & u^i \le b^i,\ \forall i & (y^i) \\ & \textstyle\sum_i x^i(\vartheta) \le 1,\ \forall \vartheta & (z(\vartheta)) \\ & x^i(\vartheta) \ge 0,\ \forall i, \vartheta \end{aligned}$$

$$\begin{aligned} \min_{w,y,z} \quad & \textstyle\sum_i y^i b^i + \sum_\vartheta z(\vartheta) \\ \text{s.t.} \quad & z(\vartheta) \ge \vartheta^i f(\vartheta) w^i,\ \forall i, \vartheta \quad & (x^i(\vartheta)) \\ & y^i + w^i \ge 1,\ \forall i & (u^i) \\ & y^i, w^i, z(\vartheta) \ge 0,\ \forall i, \vartheta \end{aligned}$$

Assume we have an optimal primal-dual pair, then if for some profile of virtual values ϑ agent i is allocated with non-zero probability, i.e., $x^i(\vartheta) > 0$ then by complementary slackness we must have for all $j \neq i$:

$$\vartheta^i w^i f(\vartheta) = z(\vartheta) \ge \vartheta^j w^j f(\vartheta)$$

where the equality follows from complementary slackness and the inequality follows from feasibility. This means that $i \in \arg\max_i \vartheta^i w^i$, except that when $f(\vartheta) = 0$.[4]

Now we still need to argue that the item is always allocated in an optimal solution. We again use complementary slackness. If the item is not completely allocated for a profile ϑ we must have $z(\vartheta) = 0$ and therefore for all agents i:

$$0 = z(\vartheta) \ge \vartheta^i w^i f(\vartheta) \ge 0$$

so $\vartheta^i w^i$ must be zero except when $f(\vartheta) = 0$.

Finally observe that even though we relaxed monotonicity in the program, the complementarity constraints imply that under any tie-breaking rule the allocation is monotone. □

[4] Assuming that these properties ($x^i(\vartheta) > 0 \Longrightarrow i \in \arg\max_i \vartheta^i w^i$ and $z(\vartheta) = 0 \Longrightarrow \vartheta^i w^i = 0$) still hold when $f(\vartheta) = 0$ never changes the optimality or the feasibility of the solution.