Games and Economic Behavior 138 (2023) 161–170

Contents lists available at ScienceDirect

Games and Economic Behavior

journal homepage: www.elsevier.com/locate/geb

Note

Ready to trade? On budget-balanced efficient trade with uncertain arrival

Daniel F. Garrett

University of Essex, Wivenhoe Park, Colchester CO4 3SQ, UK

ARTICLE INFO

Article history:
Received 12 April 2022
Available online 29 December 2022

JEL classification:
D82

Keywords:
Dynamic mechanism design
Repeated trade
Budget balance
Dynamic arrivals
Participation constraints

ABSTRACT

This paper studies the design of efficient mechanisms for repeated bilateral trade in settings where (i) traders' values and costs evolve randomly with time, and (ii) the traders become ready and available to participate in the mechanism at random times. Under a weak condition, analogous to the non-overlapping supports condition of Myerson and Satterthwaite (1983), efficient trade is only feasible if the mechanism runs an expected budget deficit. The smallest such deficit is attainable by a sequence of static mechanisms.

1. Introduction

Several papers have reversed the conclusion of Myerson and Satterthwaite's (1983) Theorem by positing dynamic environments where preferences evolve stochastically with time. These include Athey and Miller (2007), Athey and Segal (2007, 2013), Skrzypacz and Toikka (2015), Yoon (2015) and Lamba (2019). The key observation is that, when trade is repeated, when players are sufficiently patient, and when values and costs are not too persistent, trade surplus that is expected in the future can be promised to players as a reward for participation, thus relaxing participation constraints. Therefore, unlike in Myerson and Satterthwaite, budget-balanced efficient trade can be possible.[1]

While the relevant private information in the above work is the trading partners' values and costs, and their evolution over time, this paper introduces an additional source of information. Namely, we suppose each party is privately informed of the date at which they become "ready to trade". We take this to be the first date at which they are able to participate in the mechanism. We argue that the additional source of information restores the impossibility of efficient trade in budget-balanced mechanisms satisfying requisite incentive and participation constraints. Under weak conditions, efficient trade necessarily leads the mechanism to run an expected budget deficit. While earlier work has studied the interaction between stochastic valuations and dynamic arrivals in settings with profit maximization (see Deb and Said, 2015, Garrett, 2016 and 2017, Ely et al., 2017, and Bergemann and Strack, 2022), the implications for budget balance under efficient allocations are new.

E-mail address: d.garrett@essex.ac.uk.

[1] Other work exhibiting the possibility of efficient and budget-balanced trade in static settings involves departures from the assumption of risk-neutral Bayesian agents with common priors. See, in particular, Wolitzky (2016) where agents are ambiguity averse and Garratt and Pycia (2016), where agents are risk averse.

https://doi.org/10.1016/j.geb.2022.12.007
0899-8256/

One way to understand the result is in view of the rents that the agents must be granted to ensure participation. Because agents arrive at random times, an efficient mechanism must permit participation at any date. Therefore, an agent who arrives at a given date has the opportunity to delay participation and mimic later arrival. To prevent such deviations, the mechanism must provide the agent with higher rents. These additional rents turn out to be so large that, because also incentive compatibility and efficiency must be satisfied, the mechanism necessarily runs an expected budget deficit. In addition, the paper shows that the smallest feasible budget deficit is equal to that from a sequence of *static* mechanisms, each designed to ensure participation in the static mechanism alone. A message is that, at least for reducing the expected budget deficit from efficient trade, truly dynamic mechanisms can be of limited use.

1.1. Related literature

Our central departure from other papers on repeated trade with evolving preferences is set out above. We take Skrzypacz and Toikka's (2015) bilateral trade model as the key point of comparison. They provide a necessary and sufficient condition for the existence of satisfactory efficient mechanisms analogous to the condition developed for static problems by Makowski and Mezzetti (1994), Krishna and Perry (1998) and Williams (1999). Their approach makes integral use of VCG mechanisms, and we employ related arguments here.

As noted, private information on arrival times has been studied in work on profit-maximizing dynamic mechanisms with evolving preferences. Like the present paper, these works typically view "arrival" as the first date an agent is *able to participate* in a mechanism, and the possibility to secretly delay participation can be an important source of information rents. In their work on the efficient "dynamic pivot mechanism", Bergemann and Valimaki (2010) instead view dynamic arrivals as players transitioning from an "inactive state" whose participation in the mechanism is recorded from the beginning. In this view, the mechanism can contract with agents at the beginning and before their "arrival". The present analysis rules out this possibility which is important for the results. Suppose all agents begin in a "pre-arrival" state and can contract at this initial point. Then, provided they have no private information at this point, an efficient mechanism can extract all rents and avoid a budget deficit.

Related, Athey and Segal (2013, p. 2477) note that their efficient "balanced mechanism" can accommodate arrivals and exits. However, our impossibility finding indicates that their result on the possibility of "self-enforcing mechanisms" (their Proposition 4) cannot be extended to dynamic arrivals in the case of bilateral trade (at least when agents who have not arrived cannot participate). This is immediate from our result since Athey and Segal's notion of "self-enforcement" requires the satisfaction of participation constraints.

It is worth noting that some other papers with dynamic arrivals, where agents also cannot contract before arrival but where preferences do not fluctuate with time, have obtained the possibility of efficient allocations with budget balance. Gershkov and Moldovanu (2010) model dynamically arriving buyers who have heterogeneous short-lived valuations for a range of available objects. The mechanism at first instance runs a budget surplus that can be distributed back to buyers to maintain budget balance. Parkes and Singh (2003) consider a quite general framework but impose a condition on preferences that ensures a budget surplus. Loertscher et al. (2022) suppose a buyer and seller arrive in each period and find efficient and budget-balanced trade can be possible. As they note, this possibility seems related to Makowski and Mezzetti (1993) and Williams (1999) which establish corresponding results in static settings with many buyers and sellers. In light of these contributions, it is important to emphasize that the impossibility result in the present paper is specific to settings where, in the corresponding static problem, efficient allocations require a budget deficit.[2]

2. Model

Arrival of traders. We consider bilateral trade set in discrete time, with at most one unit of a perishable good sold each period. To fix ideas, and to ensure that each agent has the opportunity to engage in repeated trade irrespective of their arrival date, we suppose that the horizon is infinite. Periods are labeled $t = 1, 2, \ldots$. Agents are labeled $i \in \{B, S\}$. We term one agent the buyer ($i = B$) and the other the seller ($i = S$). These agents become "ready to trade" (equivalently, "arrive" to the market) at some dates τ_B and τ_S, which are the first dates they can enter a contract. The ability to contract and participate in the mechanism persists for the rest of time; in particular, neither buyer nor seller exit after, respectively, τ_B or τ_S.

Payoffs. In each period t, an allocation $x_t \in \{0, 1\}$ is determined with $x_t = 1$ if the seller trades the good with the buyer. For $t \geq \tau_B$, the buyer's resulting period-t payoff is $v_t x_t + p_{B,t}$, where v_t is the buyer's (private) period-t value and $p_{B,t}$ is the date-t transfer paid to the buyer. For $t \geq \tau_S$, the seller's period-t payoff is $p_{S,t} - c_t x_t$, where c_t is the seller's (private) cost and $p_{S,t}$ the transfer paid to the seller. Both agents have a common discount factor $\delta \in (0, 1)$. Throughout, we refer to v_t

[2] One of our findings is more general, however: in seeking to maximize revenue subject to efficiency, we can typically do no better than a sequence of static mechanisms, each of which induces agent participation when run on its own. In particular, the arguments in this paper can be applied to settings like repeated auctions with many buyers or repeated trade with many buyers and many sellers; however, these same arguments need not imply the impossibility of budget-balanced efficient mechanisms in these cases.

and c_t as the "payoff types" of the buyer and seller respectively, to distinguish from private information on the arrival times τ_B and τ_S. We denote vectors of payoff types by $v_s^t = (v_s, v_{s+1}, \ldots, v_t)$ and $c_s^t = (c_s, c_{s+1}, \ldots, c_t)$.

Note that the payments $p_{B,t}$ and $p_{S,t}$ can be viewed as coming from a third-party "broker" who has no private information. The focus will be on whether this broker can avoid an ex-ante budget deficit.

Stochastic processes. Each agent $i \in \{B, S\}$ independently draws an arrival time τ_i from a distribution H_i with full support on the set of periods $\mathbb{N}$. As noted, date τ_i is the first date that agent i can participate in (equivalently, communicate with) the mechanism, and is i's private information. Below, we will abuse notation by writing $v_t = \emptyset$ if the buyer has not arrived by time t (although $\emptyset$ is not a "payoff type"). Similarly, we write $c_t = \emptyset$ if the seller has not arrived by time t.

The evolution of payoff types is also independent across agents. The set of possible payoff types at date t for the buyer is denoted $\mathcal{V}_t = [\underline{v}_t, \bar{v}_t]$, with $0 \le \underline{v}_t < \bar{v}_t < \infty$, while for the seller it is $\mathcal{C}_t = [\underline{c}_t, \bar{c}_t]$, with $0 \le \underline{c}_t < \bar{c}_t < \infty$. Assume the sets $\cup_{t \ge 1} \mathcal{V}_t$ and $\cup_{t \ge 1} \mathcal{C}_t$ are bounded.

If the buyer arrives at time τ_B, he draws his payoff type v_{τ_B} from an absolutely continuous distribution $F_{\tau_B}^{In}$ with full support on $\mathcal{V}_{\tau_B}$.[3] Subsequently, at each date $t > \tau_B$, if his date $t-1$ payoff type is $v_{t-1} \in \mathcal{V}_{t-1}$, then he draws v_t from an absolutely continuous conditional distribution $F_t^{Tr}(v_t | v_{t-1})$ with support on an interval $[\underline{v}_t(v_{t-1}), \bar{v}_t(v_{t-1})] \subset \mathcal{V}_t$. Assume further that, for each v_t, $F_t^{Tr}(v_t | v_{t-1})$ is continuous in v_{t-1}.

Similarly, if the seller arrives at time τ_S, she draws her payoff type c_{τ_S} from an absolutely continuous function $G_{\tau_S}^{In}$ with full support on $\mathcal{C}_{\tau_S}$. Subsequently, at each date $t > \tau_S$, if her date $t-1$ payoff type is $c_{t-1} \in \mathcal{C}_{t-1}$, then she draws c_t from an absolutely continuous conditional distribution $G_t^{Tr}(c_t | c_{t-1})$ with support on an interval $[\underline{c}_t(c_{t-1}), \bar{c}_t(c_{t-1})] \subset \mathcal{C}_t$. Assume further that, for each c_t, $G_t^{Tr}(c_t | c_{t-1})$ is continuous in c_{t-1}.

The above description of the stochastic processes for payoff types encodes our assumption that types evolve according to a (possibly time-varying) first-order Markov process. The restriction that the supports of the conditional distributions be contained in the supports of the initial distributions, i.e. $\mathcal{V}_t$ and $\mathcal{C}_t$ respectively, is arguably quite mild. It is implied if we take $F_{\tau_B}^{In}$ and $G_{\tau_S}^{In}$, for $\tau_B, \tau_S \ge 2$, to be the marginal distributions at dates τ_B and τ_S of the payoff-type process conditional on arrival at date 1. In this case, we can view each agent's payoff type as following a latent process from date 1, with the arrival times τ_B and τ_S determined independently of the process.

We presently leave further restrictions on the evolution of payoff types unspecified, but will follow Skrzypacz and Toikka (2015) in requiring that a certain "payoff-equivalence property" holds. This property (introduced formally in Assumption 1 below) can be shown to hold under mild additional restrictions on the stochastic process (see Pavan et al., 2014, as well as Skrzypacz and Toikka).

Mechanisms. Without loss of generality, we study direct mechanisms. Each agent i makes a report of their payoff type on the first date of participation $\hat{\tau}_i$, and then continues to provide updates of these types at each date. The buyer's initial report is restricted to come from $\mathcal{V}_{\hat{\tau}_B}$. Then, if the buyer reported $\hat{v}_{t-1}$ at date $t-1$, he makes a report at date t restricted to be in $[\underline{v}_t(\hat{v}_{t-1}), \bar{v}_t(\hat{v}_{t-1})]$, which recall is the support of $F_t^{Tr}(\cdot | \hat{v}_{t-1})$. Similarly, the seller's initial report is restricted to come from $\mathcal{C}_{\hat{\tau}_S}$. Then, if the seller reported $\hat{c}_{t-1}$ at date $t-1$, she makes a report in $[\underline{c}_t(\hat{c}_{t-1}), \bar{c}_t(\hat{c}_{t-1})]$ at date t. The reports of the buyer and seller up to date t may then be denoted $\hat{v}_{\hat{\tau}_B}^t$ and $\hat{c}_{\hat{\tau}_S}^t$ respectively.

The mechanism's decision rule $\mu = \langle x_t, p_{B,t}, p_{S,t} \rangle_{t \ge 1}$ then specifies a sequence of allocations x_t for each date t and transfers $p_{B,t}$ and $p_{S,t}$ to the buyer and seller respectively. A date-t allocation is $x_t\left(\hat{v}_{\hat{\tau}_B}^t, \hat{c}_{\hat{\tau}_S}^t\right) \in \{0, 1\}$ for each possible pair of report sequences $\left(\hat{v}_{\hat{\tau}_B}^t, \hat{c}_{\hat{\tau}_S}^t\right)$.[4] Here, albeit abusively, if the buyer has not participated by date t we understand $\hat{v}_{\hat{\tau}_B}^t$ to equal $\emptyset$ and similarly $\hat{c}_{\hat{\tau}_B}^t = \emptyset$ if the seller has not participated. If either player has not participated then there is no trade, i.e. $x_t\left(\hat{v}_{\hat{\tau}_B}^t, \hat{c}_{\hat{\tau}_S}^t\right) = 0$. At any date t and for reports $\left(\hat{v}_{\hat{\tau}_B}^t, \hat{c}_{\hat{\tau}_S}^t\right)$, payments are $p_{Bt}\left(\hat{v}_{\hat{\tau}_B}^t, \hat{c}_{\hat{\tau}_S}^t\right) \in \mathbb{R}$ to the buyer and $p_{St}\left(\hat{c}_{\hat{\tau}_S}^t, \hat{v}_{\hat{\tau}_B}^t\right) \in \mathbb{R}$ to the seller. We consider payments to any player yet to participate in the mechanism to be zero.

Throughout, transfers are assumed to be measurable functions of the agents' reports. Transfers are also required to be bounded; we refer to this property of mechanisms as *bounded transfers* (BT). Boundedness will ensure that the NPV of expected total payments by the broker is well-defined in any feasible mechanism.

It may now be helpful to delineate the timing of events in each period. At the beginning of each period t, the buyer has made a sequence of reports $\hat{v}_{\hat{\tau}_B}^{t-1} \in \Pi_{s=\hat{\tau}_B}^{t-1} \mathcal{V}_s$ if he already participated at $\hat{\tau}_B < t$. If he arrived before date t, then he draws a date-t payoff type v_t from the distribution $F_t^{Tr}(v_t | v_{t-1})$, where v_{t-1} denotes his true type at $t-1$. If he arrives at date t, then he draws v_t from F_t^{In}. Having drawn his date-t valuation, he then makes a report $\hat{v}_t$ to the mechanism. If this is his first report, then it amounts to a claim that his arrival time is $\hat{\tau}_B = t$. Recall that the buyer cannot report to the mechanism if he has not yet arrived.

[3] Note that, if we think of the buyer as having a latent valuation that evolves from the first period, then the distribution $F_{\tau_B}^{In}$ is the distribution of valuations conditional on arrival at date τ_B. This distribution could reflect correlation between the buyer's value and his arrival time; e.g., if the buyer were more likely to arrive when his value is high. The analogous comments are applicable to the seller.

[4] Note the notational convention that the length of the report sequences that are arguments to the payment and allocation rules indicate the arrival times of the agents.

The timing of events for the seller is symmetric and omitted. The buyer and seller can be thought of as making any reports simultaneously within each period. Given the buyer and seller reports up to date t, the allocation $x_t\left(\hat{v}^t_{\hat{\tau}_B}, \hat{c}^t_{\hat{\tau}_S}\right)$ is determined and the transfers $p_{B,t}\left(\hat{v}^t_{\hat{\tau}_B}, \hat{c}^t_{\hat{\tau}_S}\right)$ and $p_{S,t}\left(\hat{c}^t_{\hat{\tau}_S}, \hat{v}^t_{\hat{\tau}_B}\right)$ are paid.

3. Analysis of satisfactory mechanisms

3.1. Preliminaries

Information. An important consideration is the amount of information available to each agent i at each date t regarding the past reports of the other agent. As Myerson (1986) noted, incentive constraints are most easily satisfied when agents are *least* informed. Since we are focused on the negative result — the impossibility of efficient trade without a budget deficit — we assume that agents remain completely uninformed about the other agent's reports. Such mechanisms have often been referred to as "blind". While we do not formally describe such environments, our negative result extends to settings where the players have more information than in the blind mechanism.[5]

Agent continuation payoffs. Consider then the blind mechanism that we identify with the decision rule μ. The buyer's expected continuation payoff in μ when reporting truthfully after a history of reports $v^{t-1}_{\tau_B}$ (if any), when his date-t value is v_t, is[6]

$$V^{\mu}_{B,t}\left(v^t_{\tau_B}\right) = \mathbb{E}\left[\sum_{s=t}^{\infty}\delta^{s-t}\left(\begin{array}{c} p_{B,s}\left(\tilde{v}^s_{\tau_B}, \tilde{c}^s_{\tilde{\tau}_S}\right) \\ +\tilde{v}_s\, x_s\left(\tilde{v}^s_{\tau_B}, \tilde{c}^s_{\tilde{\tau}_S}\right)\end{array}\right)\middle| \tilde{v}^t_{\tilde{\tau}_B} = v^t_{\tau_B}\right].$$

Analogously, the seller's expected continuation payoff in μ when reporting truthfully after a history of reports $c^{t-1}_{\tau_S}$ (if any), when her date-t cost is c_t, is

$$V^{\mu}_{S,t}\left(c^t_{\tau_S}\right) = \mathbb{E}\left[\sum_{s=t}^{\infty}\delta^{s-t}\left(\begin{array}{c} p_{S,s}\left(\tilde{c}^s_{\tau_S}, \tilde{v}^s_{\tilde{\tau}_B}\right) \\ -\tilde{c}_s\, x_s\left(\tilde{v}^s_{\tilde{\tau}_B}, \tilde{c}^s_{\tau_S}\right)\end{array}\right)\middle| \tilde{c}^t_{\tilde{\tau}_S} = c^t_{\tau_S}\right]$$

Incentive compatibility. To define incentive compatibility, suppose that each agent i reports to the mechanism on the arrival date τ_i. We then say that our blind mechanism is *Bayesian incentive compatible* (*BIC*) if, for each player i and arrival date τ_i, expected payoffs are maximized by reporting payoff types truthfully (assuming participation on the arrival date τ_i). Thus, for the buyer, $V^{\mu}_{B,\tau_B}\left(v_{\tau_B}\right)$ is equal to the supremum of his expected continuation payoff over all possible reporting strategies, given participation at the arrival date τ_B, and given an initial valuation $v_{\tau_B} \in \mathcal{V}_{\tau_B}$. The analogous statement holds for the seller.

Participation constraints. We assume that agents can commit to their future participation in the mechanism. Hence, we impose a sequence of constraints to ensure agents participate upon arrival, but there will be no constraints relating to continued participation.[7] Clearly, this only strengthens our results regarding the absence of satisfactory mechanisms for efficient trade.

The fact that each agent i's payoff type is a sufficient statistic for the evolution of their future types simplifies the participation constraints. Take the example of the buyer. Suppose he arrives at a given date t' and delays participation until a date $t'' > t'$, drawing a value $v_{t''}$ at that date. Then, given the restriction on the supports of $\left(F^{Tr}_t\right)_{t\geq 2}$, we note $v_{t''} \in \mathcal{V}_{t''}$. This means that he is in the same position at date t'' as if he arrived on that date with initial valuation $v_{t''}$, and can expect the same discounted payoff. The analogous statements hold for the seller.

Given the above, we specify an agent strategy of participating at all dates such that participation has not yet occurred, irrespective of the realization of the actual arrival time (or past information). Since transfers are bounded, per-period payoffs are bounded, and together with $\delta < 1$ this permits application of the one-shot deviation principle. It is then enough to check deviations in which each agent delays participation by one period. Because the process for valuations is Markov, the buyer has no profitable one-shot deviations if, at all τ_B and for all $v_{\tau_B} \in \mathcal{V}_{\tau_B}$,

$$V^{\mu}_{B,\tau_B}\left(v_{\tau_B}\right) \geq \delta\mathbb{E}\left[V^{\mu}_{B,\tau_B+1}\left(\tilde{v}_{\tau_B+1}\right)|\tilde{v}_{\tau_B} = v_{\tau_B}\right] \quad (1)$$

The seller has no profitable one-shot deviations if, for all τ_S and for all $c_{\tau_S} \in \mathcal{C}_{\tau_S}$,

[5] Note that among mechanisms minimizing the broker's expected deficit, it can be of interest to investigate those mechanisms which are fully public, where each agent learns the past reports of the other. We considered this possibility in an earlier version of the paper.

[6] We use tildes to denote random variables.

[7] By the revelation principle, it is without loss of generality to restrict to mechanisms where agents participate upon arrival.

$$V^{\mu}_{S,\tau_S}\left(c_{\tau_S}\right) \geq \delta \mathbb{E}\left[V^{\mu}_{S,\tau_S+1}\left(\tilde{c}_{\tau_S+1}\right)|\tilde{c}_{\tau_S}=c_{\tau_S}\right]. \tag{2}$$

We say that a mechanism in which all these constraints are satisfied satisfies "blind participation constraints" (or *BPC*).

It is important to note that these participation constraints will imply that agent rents accumulate in the following sense. Ensuring participation of player i at date τ_i requires satisfaction of one of the participation constraints (1) or (2). This requires enough rents be paid for participation at τ_i *relative* to what is available if participation is delayed to τ_i+1. However, rents at τ_i+1 must also be elevated to deter delay by an agent who is yet to participate at τ_i+1. This contributes to the rents that must be available for participation at date τ_i. Using this argument, increased rents that must be made available to guarantee participation from an arrival date far in the future lead to increased rents needed to guarantee participation also at the beginning. This (backwards) accumulation of rents will be the key to understanding our impossibility result below.

Efficient mechanisms. Our focus is on mechanisms that are *efficient* (E), i.e. those which set, for all t and all $v^t_{\tau_B}$ and $c^t_{\tau_S}$, $x_t\left(v^t_{\tau_B}, c^t_{\tau_S}\right) = x^E\left(v_t, c_t\right)$, the efficient allocation which is taken to equal one in case $t \geq \max\{\tau_B, \tau_S\}$ and $v_t \geq c_t$, and zero otherwise.

Budget balance. We say that a mechanism μ satisfies *budget balance* (BB) if two requirements are met. First, payments are zero unless both players have arrived. Second, for all arrival times τ_B and τ_S, for all dates $t \geq \max\{\tau_B, \tau_S\}$, all feasible sequences of values and costs $v^t_{\tau_B}$ and $c^t_{\tau_S}$, we have $p_{B,t}\left(v^t_{\tau_B}, c^t_{\tau_S}\right) + p_{S,t}\left(v^t_{\tau_B}, c^t_{\tau_S}\right) = 0$.

Payoff equivalence. Following Skrzypacz and Toikka (2015), we restrict attention to environments in which a version of "payoff equivalence" holds. We assume the following.

Assumption 1. The stochastic processes for payoff types satisfy the "payoff-equivalence property" meaning that the following holds. Consider any two BIC blind mechanisms $\mu = \langle x_t, p_{B,t}, p_{S,t}\rangle_{t\geq 1}$ and $\mu' = \left\langle x'_t, p'_{B,t}, p'_{S,t}\right\rangle_{t\geq 1}$ satisfying $x_t = x'_t$ for all t. There exist real-valued scalars $\left(b_{B,\tau_B}, b_{S,\tau_S}\right)_{\tau_B,\tau_S\geq 1}$ such that (i) for each date τ_B and each $v_{\tau_B} \in \mathcal{V}_{\tau_B}$, $V^{\mu}_{B,\tau_B}\left(v_{\tau_B}\right) = b_{B,\tau_B} + V^{\mu'}_{B,\tau_B}\left(v_{\tau_B}\right)$; and (ii) for each date τ_S and each $c_{\tau_S} \in \mathcal{C}_{\tau_S}$, $V^{\mu}_{S,\tau_S}\left(c_{\tau_S}\right) = b_{S,\tau_S} + V^{\mu'}_{S,\tau_S}\left(c_{\tau_S}\right)$. Moreover, for any BIC blind mechanism μ, any initial participation dates τ_B and τ_S, $V^{\mu}_{B,\tau_B}(\cdot)$ and $V^{\mu}_{S,\tau_S}(\cdot)$ are absolutely continuous.

Considering blind mechanisms, the relevant notion of payoff equivalence is simply that each agent i's expected payoff from participating at date τ_i is the same, up to a constant b_{i,τ_i}, for any BIC mechanism with the same allocation rule. Sufficient conditions for our assumption to hold can be found in Pavan et al. (2014), where they consider a dynamic envelope theorem which also yields the absolute continuity of payoffs that we assume.[8]

3.2. Main results

With these definitions in hand, we can now state the first step of the analysis.

Lemma 3.1. *If a blind mechanism μ maximizes the broker's expected surplus (equivalently, minimizes the broker's expected deficit) among mechanisms satisfying E, BIC and BPC then the following are true. For all τ_B,*

$$\inf_{v_{\tau_B}\in\mathcal{V}_{\tau_B}} \left\{V^{\mu}_{B,\tau_B}\left(v_{\tau_B}\right) - \delta\mathbb{E}\left[V^{\mu}_{B,\tau_B+1}\left(\tilde{v}_{\tau_B+1}\right) \mid \tilde{v}_{\tau_B} = v_{\tau_B}\right]\right\} = 0 \tag{3}$$

and for all τ_S,

$$\inf_{c_{\tau_S}\in\mathcal{C}_{\tau_S}} \left\{V^{\mu}_{S,\tau_S}\left(c_{\tau_S}\right) - \delta\mathbb{E}\left[V^{\mu}_{S,\tau_S+1}\left(\tilde{c}_{\tau_S+1}\right) \mid \tilde{c}_{\tau_S} = c_{\tau_S}\right]\right\} = 0. \tag{4}$$

To understand this result consider the case of the buyer. In order for immediate participation in the mechanism to always be optimal, there must be no profitable one-shot deviations, which means that (1) holds at all dates τ_B, and for all $v_{\tau_B} \in \mathcal{V}_{\tau_B}$. For any date τ_B at which Equation (3) fails to hold, we can reduce the payment to the buyer at date τ_B by a constant that does not depend on his payoff type v_{τ_B}. This increases the broker's expected surplus (or decreases the deficit), and, for a small enough change, it leaves all of the constraints (1) intact. Specifically, it tightens the slack constraint at τ_B, while for $\tau_B > 1$, it relaxes the constraint at $\tau_B - 1$ by reducing the buyer's option value of delay.

We now show that, to maximize the broker's surplus, it is enough to rely on a sequence of static mechanisms, with the understanding that each static mechanism elicits participation from any agent who has already arrived. We begin by considering the static VCG mechanism which ensures each agent a payoff equal to the surplus from trade, i.e. $(v_t - c_t)$

[8] The conditions for their result essentially ensure that agent payoffs are not "too sensitive" to small changes in type. This includes ruling out the future evolution of types being too sensitive to small changes in type, as formalized by the notion of impulse responses.

$x^E(v_t, c_t)$ at date t. Note that a fixed fee can be charged to the buyer and seller while still ensuring a non-negative payoff from participation. The buyer's date-t fee can be positive if type $\underline{v}_t$ expects a positive surplus from trade; similarly, the seller's date-t fee can be positive if type $\bar{c}_t$ expects a positive surplus from trade. We therefore define a "VCG-plus mechanism" to be one with the efficient allocation rule x^E and with payments in case both agents have arrived by date t (i.e., $\tau_B, \tau_S \leq t$) given by

$$p_{B,t}^{VCG-plus}\left(v_{\tau_B}^t, c_{\tau_S}^t\right) = -c_t\, x^E\left(v_t, c_t\right) - \mathbb{E}\left[\left(\underline{v}_t - \tilde{c}_t\right)\, x^E\left(\underline{v}_t, \tilde{c}_t\right) \mid \tilde{c}_t \neq \emptyset\right] \tag{5}$$

for the buyer, and

$$p_{S,t}^{VCG-plus}\left(v_{\tau_B}^t, c_{\tau_S}^t\right) = v_t\, x^E\left(v_t, c_t\right) - \mathbb{E}\left[\left(\tilde{v}_t - \bar{c}_t\right)\, x^E\left(\tilde{v}_t, \bar{c}_t\right) \mid \tilde{v}_t \neq \emptyset\right] \tag{6}$$

for the seller. If one or both agents have not arrived by date t, then there is no trade and no transfers are made. The VCG-plus mechanism may be viewed as a sequence of static mechanisms that elicit participation from any player that has already arrived at each date. However, in the discussion below, we will view this now as a dynamic mechanism, consistent with other such mechanisms in the paper, by requiring that an agent who participates in the mechanism is compelled to continue. We find the following.

Proposition 3.1. *The broker's expected surplus is maximized (among mechanisms satisfying E, BIC and BPC) by running a VCG-plus mechanism. That is, it is maximized by setting* $\mu = \left(x^E, p_{B,t}^{VCG-plus}, p_{S,t}^{VCG-plus}\right)_{t\geq 1}$.

The reason for this result is as follows. Running a sequence of VCG mechanisms, one each period, is a simple way to implement the efficient allocation. Now suppose that transfers are augmented by fixed (i.e., type-independent) participation fees as in (5) and (6), obtaining the VCG-plus mechanism described in the proposition. Then, at each date t, the buyer with the lowest value $\underline{v}_t$ and the seller with the highest cost $\bar{c}_t$ obtain an additional expected payoff of zero by participating at date t, rather than delaying participation to the subsequent period. This shows that the optimality conditions in Lemma 3.1 are satisfied for both buyer and seller at each date t. Using the "payoff equivalence property" of Assumption 1, we then show that the VCG-plus mechanism generates the same expected surplus for the broker as any other mechanism satisfying E, BIC and BPC, and the conditions of Lemma 3.1. That is, the VCG-plus mechanism is optimal for the broker in implementing efficient trade.

Proposition 3.1 shows that the broker's net revenue can be maximized by running a sequence of static mechanisms where, if each mechanism were run on its own, the buyer and seller would want to participate. In this sense, there is no way for the broker to benefit from repeated interactions to limit agent rents. It is worth comparing this finding to broker-optimal mechanisms in settings without dynamic arrivals, where the buyer and seller are present from the beginning, as for instance in Skrzypacz and Toikka (2015). There, if participation constraints are only imposed at the initial date (what Skrzypacz and Toikka call IR_0), maximal broker revenue can be obtained by charging up-front participation fees for the right to participate in a sequence of efficient static mechanisms. These payments could also be spread over time in different ways. What is generally *not possible* in the environment without dynamic arrivals, however, is that the broker's maximal revenue be obtained by a sequence of static mechanisms that, each run on its own, would elicit buyer and seller participation (i.e., the broker can do better than this).

What is the broker's expected surplus in the broker-optimal efficient mechanism? First, recall that the VCG mechanisms that we took as our starting point imply a budget deficit in the bilateral trade problem equal to the total surplus $(v_t - c_t)$ $x^E(v_t, c_t)$ at each date t. However, this has to be balanced against the fees collected in the VCG-plus mechanism. We have the following result.

Proposition 3.2. *The largest value of the broker's expected surplus in a blind mechanism satisfying E, BIC and BPC is*

$$R = \sum_{t=1}^{\infty} \delta^{t-1} H_B(t)\, H_S(t)\, \Psi_t, \tag{7}$$

where

$$\Psi_t = -\mathbb{E}\left[\left(\tilde{v}_t - \tilde{c}_t\right) x^E\left(\tilde{v}_t, \tilde{c}_t\right) | \tilde{v}_t \neq \emptyset, \tilde{c}_t \neq \emptyset\right] + \mathbb{E}\left[\left(\underline{v}_t - \tilde{c}_t\right)\, x^E\left(\underline{v}_t, \tilde{c}_t\right) | \tilde{c}_t \neq \emptyset\right] + \mathbb{E}\left[\left(\tilde{v}_t - \bar{c}_t\right)\, x^E\left(\tilde{v}_t, \bar{c}_t\right) | \tilde{v}_t \neq \emptyset\right]. \tag{8}$$

There exists a blind mechanism that satisfies BB, E, BIC and BPC if and only if $R \geq 0$. A sufficient condition for $R < 0$ is that $\mathcal{V}_t \cap \mathcal{C}_t$ has positive length for each t, meaning that the distributions $\Pr\left(\tilde{v}_t \leq v_t | \tilde{v}_{B,t} \neq \emptyset\right)$ and $\Pr\left(\tilde{c}_t \leq c_t | \tilde{c}_t \neq \emptyset\right)$ have "overlapping supports" for all t.

Equation (7) should be understood as a weighted average of the broker's maximal expected surplus under efficient trade in static mechanisms, as calculated, for instance, by Makowski and Mezzetti (1994). In particular, if we consider the broker-optimal (efficient) static mechanism at some date t, with type distributions determined by conditioning only on the arrival of both agents by date t, then the broker's expected surplus is precisely Ψ_t. The weights in (7) comprise the discount factor δ^{t-1} and the probability that both agents arrive by date t, $H_B(t) H_S(t)$.[9]

If $\Psi_t < 0$ for all t, as is the case when the supports of the aforementioned distributions overlap (a result originally due to Myerson and Satterthwaite, 1983), then budget-balanced efficient trade is infeasible. Conversely, if $\Psi_t \geq 0$ at a given date t, then an efficient and budget-balanced static mechanism exists at date t; in this case efficient trade in the static mechanism can be implemented through a type-independent posted price. Hence, if $\Psi_t \geq 0$ for all t, budget-balanced trade is achievable through a sequence of posted prices.

If $\Psi_t < 0$ for some t and yet the expression in (7) is non-negative, then a blind mechanism can be chosen to satisfy E, BB, BIC and BPC. The mechanism described in the Appendix is simply a sequence of (static) AGV mechanisms with additional fixed (i.e., independent of payoff type) payments between the agents. Note that, given that $\Psi_t < 0$ for some t, it can be necessary for some of these payments to be made after the play of a given static mechanism in the sequence. In other words, playing the mechanism at a given date t gives rise to dynamic obligations, and these obligations are essential for "spreading" surplus across players with different arrival times, thus ensuring willingness to participate at each date. In this sense, while a sequence of static mechanisms is enough to maximize the broker's expected surplus R, truly "dynamic" mechanisms may be needed to obtain budget balance. It is worth commenting that the "spreading" of surplus across time that is needed here is quite different from the "rebalancing" of incentive payments found in Athey and Segal (2013), where payments need to be carefully designed to ensure truthful reporting of valuations. It is perhaps more closely related to the idea discussed in Section 4.2 of Gershkov and Moldovanu (2010) on "offline mechanisms" where payments by one buyer are redistributed to later arrivals. The point is that the additional payments needed between players affect the levels of expected payoffs, but do not affect incentives for the truthful revelation of values and costs.

4. Conclusions

We conclude by reiterating a central theme of this paper. Following Myerson and Satterthwaite (1983), among others, it is well understood that balanced-budget requirements provide a severe impediment for efficient trade under standard incentive-compatibility and participation constraints. Following a number of elegant contributions, repeated trade, with dynamically evolving payoff types, has since emerged as a way to restore efficient allocations. A key message of the present paper, then, is that such a conclusion can be too optimistic. If agents have private information on their readiness to participate in a dynamic mechanism (i.e., if they arrive over time and their arrival dates are privately known), then efficient budget-neutral trade can be (much) more difficult to sustain. As in the classic bilateral trade problem, overlapping supports for buyer and seller values is enough to render budget-balanced and efficient trade infeasible.

Data availability

No data was used for the research described in the article.

Acknowledgments

Thanks to two anonymous referees and an Advisory Editor for helpful comments on the paper. This paper has also benefited from comments from Alessandro Bonatti, Bruno Jullien, Rohit Lamba, Niccolo Lomys, Alessandro Pavan, and Michael Whinston. Thanks to Zudik Hernández Gomis for research assistance. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 714147). The initial phases of this project took place at Toulouse School of Economics, where it was supported by the French National Research Agency (ANR) under the Investments for the Future (Investissements d'Avenir) program (grant ANR-17-EURE-0010). Any error is my own.

Appendix A. Proofs of results

Proof of Lemma 3.1. Suppose that one of the equalities in the lemma fails to hold. Consider that for the buyer at a date τ_B, since the case for the seller is analogous. Given our assumption on the evolution of payoff types and focus on BIC mechanisms, a necessary condition for the buyer to be willing to participate at date τ_B when his type is v_{τ_B} is given by (1). This implies that the expression in the left-hand side of (3) is non-negative in any blind mechanism satisfying BIC and BPC. Therefore, at τ_B the expression must be strictly positive, and so we can increase the broker's surplus simply by reducing

[9] Note that the broker's expected surplus (7) thus depends only on the distributions of arrival times and the marginal distributions of payoff types, $\Pr\left(\tilde{\theta}_{B,t} \leq \theta_{B,t} | \tilde{\theta}_{B,t} \neq \emptyset\right)$ and $\Pr\left(\tilde{\theta}_{S,t} \leq \theta_{S,t} | \tilde{\theta}_{S,t} \neq \emptyset\right)$, at each date t. Hence, unlike the existing literature on repeated trade reviewed above, the degree of persistence of agent payoff types over time is irrelevant for calculating the broker's surplus.

$p_{B,\tau_B}\left(v_{\tau_B}, c_{\tau_S}^{\tau_B}\right)$ by a uniform constant less than the positive value in question. This reduces $V_{B,\tau_B}^{\mu}\left(v_{\tau_B}\right)$ by the same amount. If $\tau_B = 1$, this yields a BIC mechanism that satisfies all the participation constraints (1). The same is true if $\tau_B > 1$, noting now that the constraint (1) at date $\tau_B - 1$ is relaxed. Note also that the buyer's transfers remain bounded, so we have found a more profitable mechanism satisfying BIC and BPC, and which still satisfies BT. □

Proof of Proposition 3.1. Consider any blind mechanism μ satisfying E, BIC, BPC and the conditions of Lemma 3.1 for all dates. Given the continuity of payoffs property in Assumption 1 and the continuity assumption on F_t^{Tr}, there exists for each t a $v_t^* \in \mathcal{V}_t$ satisfying

$$V_{B,t}^{\mu}\left(v_t^*\right) = \delta\left[V_{B,t+1}^{\mu}\left(\tilde{v}_{t+1}\right)|\tilde{v}_t = v_t^*\right].$$

Similarly, for each t, there is a c_t^* satisfying

$$V_{S,t}^{\mu}\left(c_t^*\right) = \delta\left[V_{S,t+1}^{\mu}\left(\tilde{c}_{t+1}\right)|\tilde{c}_t = c_t^*\right].$$

Further, by the payoff equivalence property in Assumption 1, for each t, we can take the values of v_t^* and c_t^* to be the same for any blind mechanism μ satisfying E, BIC, BPC and the conditions of Lemma 3.1.

For each t,

$$V_{B,t}^{\mu}\left(v_t^*\right) = \delta V_{B,t+1}^{\mu}\left(v_{t+1}^*\right) + \delta\mathbb{E}\left[\left(V_{B,t+1}^{\mu}\left(\tilde{v}_{t+1}\right) - V_{B,t+1}^{\mu}\left(v_{t+1}^*\right)\right)|\tilde{v}_t = v_t^*\right].$$

Iterating, we can write, for any $l \in \mathbb{N}$,

$$V_{B,t}^{\mu}\left(v_t^*\right) = \sum_{s=1}^{l}\delta^s\mathbb{E}\left[\left(V_{B,t+s}^{\mu}\left(\tilde{v}_{t+s}\right) - V_{B,t+s}^{\mu}\left(v_{t+s}^*\right)\right)|\tilde{v}_{t+s-1} = v_{t+s-1}^*\right] + \delta^l V_{B,t+l}^{\mu}\left(v_{t+l}^*\right).$$

Because payoff types and transfers are bounded (by BT), $V_{B,t+l}^{\mu}\left(v_{t+l}^*\right)$ is bounded across l. We can conclude that $\lim_{l\to+\infty}\delta^l V_{B,t+l}^{\mu}\left(v_{t+l}^*\right) = 0$. Therefore, for a mechanism μ satisfying E, BIC, BPC and the conditions of Lemma 3.1 for all dates, we have for all t,

$$V_{B,t}^{\mu}\left(v_t^*\right) = \sum_{s=1}^{\infty}\delta^s\mathbb{E}\left[\left(V_{B,t+s}^{\mu}\left(\tilde{v}_{t+s}\right) - V_{B,t+s}^{\mu}\left(v_{t+s}^*\right)\right)|\tilde{v}_{t+s-1} = v_{t+s-1}^*\right].$$

Similarly, we can conclude that, for all t,

$$V_{S,t}^{\mu}\left(c_t^*\right) = \sum_{s=1}^{\infty}\delta^s\mathbb{E}\left[\left(V_{S,t+s}^{\mu}\left(\tilde{c}_{t+s}\right) - V_{S,t+s}^{\mu}\left(c_{t+s}^*\right)\right)|\tilde{c}_{t+s-1} = c_{t+s-1}^*\right].$$

By the payoff equivalence property of Assumption 1, for all t and s, $V_{B,t+s}^{\mu}\left(v_{t+s}\right) - V_{B,t+s}^{\mu}\left(v_{t+s}^*\right)$ is the same across all blind mechanisms μ satisfying BIC and E. Similarly, $V_{S,t+s}^{\mu}\left(c_{t+s}\right) - V_{S,t+s}^{\mu}\left(c_{t+s}^*\right)$ is the same across all blind mechanisms μ satisfying BIC and E. Therefore, buyer and seller payoffs V_{B,τ_B}^{μ} and V_{S,τ_S}^{μ} are uniquely determined in any efficient mechanism μ satisfying the conditions in Lemma 3.1 at all dates. These values pin down the broker's expected net revenue, as this is the expected surplus from trade less expected agent rents; that is,

$$\mathbb{E}\left[\sum_{t=\max\{\tilde{\tau}_B,\tilde{\tau}_S\}}^{\infty}\delta^{t-1}\left(\tilde{v}_t - \tilde{c}_t\right)x^E\left(\tilde{v}_t, \tilde{c}_t\right)\right] - \mathbb{E}\left[\delta^{\tilde{\tau}_B-1}V_{B,\tilde{\tau}_B}^{\mu}\left(\tilde{v}_{\tilde{\tau}_B}\right)\right] - \mathbb{E}\left[\delta^{\tilde{\tau}_S-1}V_{S,\tilde{\tau}_S}^{\mu}\left(\tilde{c}_{\tilde{\tau}_S}\right)\right].$$

One mechanism that attains these profits is the VCG-plus mechanism specified in the proposition. Such mechanisms are BIC because each static VCG mechanism is incentive compatible, and because the future payment rule faced by the agents does not depend on current reports. To see that the mechanism satisfies BPC, note that any agent who has not yet participated in the VCG-plus mechanism has a weakly higher payoff from participating in the mechanism immediately rather than with a one-period delay. This follows from the choice of fixed payments appended to the static VCG mechanisms as described in the main text. The buyer with the lowest value and the seller with the highest cost at a given date t are precisely indifferent between participating at t and waiting and participating at date $t+1$ (i.e., $v_t^* = \underline{v}_t$ and $c_t^* = \bar{c}_t$). Hence, the conditions of Lemma 3.1 are satisfied. It follows from the above that the buyer and seller rents are equal to those in any other mechanism satisfying E, BIC, BPC and the conditions in Lemma 3.1, establishing that the mechanism is optimal among efficient mechanisms. □

Proof of Proposition 3.2. The expression (7) for R follows from the arguments in the main text. As we have noted, budget-balanced efficient trade is infeasible if $R<0$. (The sufficient condition for $R<0$ comes from Myerson and Satterthwaite, 1983.)

If $R\geq 0$, then we can construct a blind mechanism satisfying BB, E, BIC and BPC as follows. First we consider static mechanisms run at each date t, assuming participation by the players provided they have arrived by that date. A date-t budget-balanced, efficient and incentive-compatible static mechanism $\mu_t^{static}=\left\langle x^E, p_{B,t}^{static}, p_{S,t}^{static}\right\rangle$ is such that

$$W_{B,t}^{\mu_t^{static}}\left(\underline{v}_t\right)+W_{S,t}^{\mu_t^{static}}\left(\bar{c}_t\right)=\Psi_t, \tag{9}$$

with Ψ_t given by (8), where (i) $W_{B,t}^{\mu_t^{static}}\left(v_t\right)$ gives the buyer's expected payoff from truthful reporting in μ_t^{static}, conditional on both agents having arrived by date t, and on the buyer's date-t value being v_t; and (ii) $W_{S,t}^{\mu_t^{static}}\left(c_t\right)$ gives the seller's expected payoff from truthful reporting in μ_t^{static}, conditional on both agents having arrived by date t, and on the seller's date-t cost being c_t. That a budget-balanced mechanism exists follows by d'Aspremont and Gérard-Varet (1979).

Equation (9) follows from a familiar argument. Let μ_t^{VCG} be a VCG mechanism offered at date t that yields both players a payoff equal to the surplus from trade, and let $W_{B,t}^{\mu_t^{VCG}}\left(v_t\right)$ be the buyer's expected payoff conditional on both agents having arrived by date t, and the buyer having date-t value v_t. Similarly, let $W_{S,t}^{\mu_t^{VCG}}\left(c_t\right)$ be the seller's expected payoff conditional on both agents having arrived by date t, and the seller having date-t value c_t. By a standard envelope argument, there are scalars a_t and b_t such that, for any v_t and c_t, we have $W_{B,t}^{\mu_t^{static}}\left(v_t\right)=W_{B,t}^{\mu_t^{VCG}}\left(v_t\right)+a_t$ and $W_{S,t}^{\mu_t^{static}}\left(c_t\right)=W_{S,t}^{\mu_t^{VCG}}\left(v_t\right)+b_t$. Note that, because μ_t^{static} is budget balanced, we have

$$\mathbb{E}\left[W_{B,t}^{\mu_t^{VCG}}\left(\tilde{v}_t\right)+a_t+W_{S,t}^{\mu_t^{VCG}}\left(\tilde{c}_t\right)+b_t|\tilde{v}_t\neq\emptyset,\tilde{c}_t\neq\emptyset\right]=\mathbb{E}\left[\left(\tilde{v}_t-\tilde{c}_t\right)x^E\left(\tilde{v}_t,\tilde{c}_t\right)|\tilde{v}_t\neq\emptyset,\tilde{c}_t\neq\emptyset\right].$$

Because the VCG mechanism in question yields each player the surplus from trade, we conclude

$$a_t+b_t=-\mathbb{E}\left[\left(\tilde{v}_t-\tilde{c}_t\right)x^E\left(\tilde{v}_t,\tilde{c}_t\right)|\tilde{v}_t\neq\emptyset,\tilde{c}_t\neq\emptyset\right].$$

Moreover,

$$\begin{aligned}W_{B,t}^{\mu_t^{static}}\left(\underline{v}_t\right)+W_{S,t}^{\mu_t^{static}}\left(\bar{c}_t\right)&=W_{B,t}^{\mu_t^{VCG}}\left(\underline{v}_t\right)+W_{S,t}^{\mu_t^{VCG}}\left(\bar{c}_t\right)+a_t+b_t\\&=\Psi_t\end{aligned}$$

as desired.

Given a sequence of budget-balanced mechanisms (say AGV mechanisms), it is then possible to construct a sequence of budget-balanced static mechanisms $\mu_t^{\#}$ such that expected payoffs from participating, with value $\underline{v}_t$ for the buyer and $\bar{c}_t$ for the seller (and assuming the other agent participates), are given by

$$\begin{aligned}&W_{B,t}^{\mu_t^{\#}}\left(\underline{v}_t\right)=\min\{0,\Psi_t\}\\&W_{S,t}^{\mu_t^{\#}}\left(\bar{c}_t\right)=\max\{0,\Psi_t\}.\end{aligned}$$

This simply requires adding "fixed" (i.e., type-independent) transfers at each date t to redistribute the expected surplus between the buyer and seller. Under our maintained assumption that each agent is blind as to whether the other is participating, the buyer's expected payoff from participating in the static mechanism $\mu_t^{\#}$ is $H_S(t)\min\{0,\Psi_t\}$ when his value for the good is $\underline{v}_t$, while the seller's is $H_B(t)\max\{0,\Psi_t\}$ when his cost is $\bar{c}_t$.

We can then modify the sequence of static mechanisms $\left(\mu_t^{\#}\right)_{t\geq 1}$ by arranging for further transfers between the agents. The result need no longer be a sequence of static mechanisms, since participation at a given date t may (at least along some realizations of uncertainty) give rise to further payments at later dates. It will be enough to ensure that each agent is willing to participate at each date for all possible payoff types, irrespective of whether he participated in the past.

One way to proceed is as follows. Consider the first date $\bar{t}_1$ such that $\Psi_{\bar{t}_1}<0$. We can require the seller when participating at dates t such that $\Psi_t>0$ to make payments to any buyer who participates in the date-$\bar{t}_1$ mechanism. For a seller participation date $t<\bar{t}_1$, require this payment to be made at date $\bar{t}_1$. The NPV of the seller's expected additional payment must be no greater than the smallest expected gain from participating in the date-t mechanism $\mu_t^{\#}$, i.e. $H_B(t)\Psi_t$. This requires that the date-$\bar{t}_1$ payment does not exceed $\frac{H_B(t)\Psi_t}{\delta^{\bar{t}_1-t}H_B\left(\bar{t}_1\right)}$. For $t>\bar{t}_1$, have the seller make the payment at date t, again so the expected payment does not exceed the smallest expected gain from participating in the date-t mechanism $\mu_t^{\#}$; that is, it must be no greater than $\frac{H_B(t)\Psi_t}{H_B\left(\bar{t}_1\right)}$. Payments may be required corresponding to multiple seller participation dates with $\Psi_t>0$, up to the point where the buyer with valuation $\underline{v}_{\bar{t}_1}$ expects zero payoff from participating in the date-$\bar{t}_1$ mechanism. We can take the date-$\bar{t}_1$ buyer to be compensated by the earliest payments possible, with payments up to the aforementioned upper bounds on seller payments until a date at which the buyer is "paid off", i.e. the buyer with valuation $\underline{v}_{\bar{t}_1}$ expects zero payoff from participating at date $\bar{t}_1$.[10] Then we proceed to the next date $\bar{t}_2$ at which $\Psi_{\bar{t}_2} < 0$. We ask the seller to make payments as before so that for seller participation $t < \bar{t}_2$, a payment is made by the seller provided $\Psi_t > 0$ and the upper bound on payments was not already met in relation to buyer participation at date $\bar{t}_1$. If $t > \bar{t}_2$, then the payment is made at date t. Total payments by the seller must be chosen so that the expected additional payments of the seller participating at each date t has NPV no greater than the expected payoff from the participation in $\mu_t^{\#}$. Again, take these payments to be as early as possible, subject to the bounds on payments by the seller. Proceeding in this way, using $R \geq 0$, it is possible to construct a sequence of budget-balanced mechanisms in which, at each date t, the "worst" types of buyer and seller, i.e. $\underline{v}_t$ and $\bar{c}_t$, are willing to participate. Then we can take the efficient dynamic mechanism μ to be given by the defined payments. (Here, we may view agents as being compelled to participate after their first participation in the mechanism, although such continued participation is by construction incentive compatible.)

It remains to check that payments remain bounded. An upper bound on the seller's additional payments to a buyer participating at date $\bar{t}$ due to the seller's participation before date $\bar{t}$ is $\frac{H_S(\bar{t})|\Psi_{\bar{t}}|}{H_S(1)}$, since with this payment the buyer then expects to receive at least $H_S(\bar{t})\,|\Psi_{\bar{t}}|$, the largest expected loss from participating in $\mu_{\bar{t}}^{\#}$. An upper bound on the seller's payment to a buyer participating at date $\bar{t}$ due to the seller's participation at a date t after date $\bar{t}$ is $\frac{H_B(t)\Psi_t}{H_B(\bar{t})}$, as mentioned above. Because payoff types are bounded, the sequence $(\Psi_s)_{s=1}^{\infty}$ is bounded, so it is easy to see that the additional payments remain bounded. Because the payments in the mechanisms $\mu_s^{\#}$, $s \geq 1$, are also bounded, we conclude that the mechanism μ defined in the previous paragraph satisfies BT. □

References

d'Aspremont, Claude, Gérard-Varet, Louis-André, 1979. Incentives and incomplete information. J. Public Econ. 11, 25–45.
Athey, Susan, Miller, David, 2007. Efficiency in repeated trade with hidden valuations. Theor. Econ. 2, 299–354.
Athey, Susan, Segal, Ilya, 2007. Designing dynamic mechanisms for dynamic bilateral trading games. Am. Econ. Rev. 97, 131–136.
Athey, Susan, Segal, Ilya, 2013. An efficient dynamic mechanism. Econometrica 81, 2463–2485.
Bergemann, Dirk, Valimaki, Juuso, 2010. The dynamic pivot mechanism. Econometrica 78, 771–789.
Bergemann, Dirk, Strack, Philipp, 2022. Progressive participation. Theor. Econ. 17, 1007–1039.
Deb, Rahul, Said, Maher, 2015. Dynamic screening with limited commitment. J. Econ. Theory 159B, 891–928.
Ely, Jeffrey, Garrett, Daniel, Hinnosaar, Toomas, 2017. Overbooking. J. Eur. Econ. Assoc. 15, 1258–1301.
Garratt, Rod, Pycia, Marek, 2016. Efficient Bilateral Trade. Mimeo. UC Santa Barbara and UCLA.
Garrett, Daniel, 2016. Intertemporal price discrimination: dynamic arrivals and changing values. Am. Econ. Rev. 106, 3275–3299.
Garrett, Daniel, 2017. Dynamic mechanism design: dynamic arrivals and changing values. Games Econ. Behav. 104, 595–612.
Gershkov, Alex, Moldovanu, Benny, 2010. Efficient sequential assignment with incomplete information. Games Econ. Behav. 68, 144–154.
Krishna, Vijay, Perry, Motty, 1998. Efficient Mechanism Design. Mimeo. Hebrew University of Jerusalem and Pennsylvania State University.
Lamba, Rohit, 2019. Efficiency with(out) Intermediation in Repeated Bilateral Trade. Mimeo. Pennsylvania State University.
Loertscher, S., Muir, E.V., Taylor, P.G., 2022. Optimal market thickness. J. Econ. Theory 200, 105383.
Makowski, Louis, Mezzetti, Claudio, 1993. The possibility of efficient mechanisms for trading an indivisible object. J. Econ. Theory 59, 451–465.
Makowski, Louis, Mezzetti, Claudio, 1994. Bayesian and weakly robust first best mechanisms: characterizations. J. Econ. Theory 64, 500–519.
Myerson, Roger, 1986. Multistage games with communication. Econometrica 54, 323–358.
Myerson, Roger, Satterthwaite, Mark, 1983. Efficient mechanisms for bilateral trade. J. Econ. Theory 29, 265–281.
Parkes, David, Singh, Satinder, 2003. An MDP-based approach to online mechanism design. In: Proceedings of the 17th Annual Conference on Neural Information Processing Systems.
Pavan, Alessandro, Segal, Ilya, Toikka, Juuso, 2014. Dynamic mechanism design: a Myersonian approach. Econometrica 82, 601–653.
Skrzypacz, Andrzej, Toikka, Juuso, 2015. Mechanisms for repeated trade. Am. Econ. J. Microecon. 7, 252–293.
Williams, Steven, 1999. A characterization of efficient, Bayesian incentive compatible mechanisms. Econ. Theory 14, 155–180.
Wolitzky, Alexander, 2016. Mechanism design with maxmin agents: theory and an application to bilateral trade. Theor. Econ. 11, 971–1004.
Yoon, Kiho, 2015. On budget balance of the dynamic pivot mechanism. Games Econ. Behav. 94, 206–213.

[10] If $R = 0$, and if $\bar{t}_1$ is the last date such that $\Psi_t < 0$, then the buyer may be paid off only in the limit. That is, it may be that the NPV of an infinite stream of payments specified from the seller exactly covers the buyer's loss in the mechanism $\mu_{\bar{t}_1}^{\#}$.