

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

Daniel F. Garrett

Toulouse School of Economics

2016

Abstract

This paper studies the design of mechanisms for repeated 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.

JEL classification: D82

Keywords: dynamic mechanism design, subsidies, budget balance, dynamic arrivals

1 Introduction

A foundational result in information economics and mechanism design is Myerson and Satterthwaite's (1983) theorem, which considers two potential traders of a single unit with relevant private information (the buyer's value and the seller's cost). The result states that any efficient, (Bayesian) incentive-compatible and (interim) individually rational mechanism runs an expected budget deficit (i.e., budget-balanced trade is impossible). Recently, however, several papers examine how the classic impossibility result can be overturned in settings with repeated trade; see Athey and Miller (2007), Athey and Segal (2007, 2013), Skrzypacz and Toikka (2015), Lamba (2013) and Yoon (2015). The key observation is that, when trade is repeated, when players are sufficiently patient, and when costs and values evolve stochastically with time, trade surplus that is expected in future can be promised to players as a reward for participation, thus relaxing participation constraints. One interpretation is that these papers call into question the universality of Myerson and Satterthwaite's impossibility result.¹

In the above work, the relevant private information is the parties' values and costs of the unit traded. The purpose of the present paper is to introduce an additional source of information: the time at which each party becomes "ready to trade". We argue that this additional source restores the impossibility of efficient trade in budget-balanced mechanisms satisfying requisite incentive constraints.

Our conclusion depends on our understanding of an agent's "readiness to trade" as the state of being ready and able to participate in a trading agreement; hence, the key friction is agents' readiness to enter a (possibly long-term) contract rather than the technological feasibility of trade per se (or the possibility that such trade is mutually beneficial).² Our notion of readiness follows from our view that agents are often either unprepared to enter a given contractual relationship, or even unaware of

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

²A common observation in the dynamic mechanism design literature has been that contracting with parties as early as possible enriches the implementable outcomes (an observation which holds, at least in a weak sense, by a revelation principle for dynamic mechanisms). For instance, inducing participation in dynamic contracts at an early stage, when parties are relatively uninformed, can permit a reduction in their information rents. This paper limits the possibilities for early contracting by assuming that parties' readiness to contract arrives stochastically over time.

the possibility, but that this may be resolved with time. In the first case, akin to the incomplete contracting literature, agents may need to devote scarce attention to understanding a (dynamic) trading agreement, attention that may become available only randomly and after a delay (see Simon, 1955, for an early discussion of the difficulties of evaluating payoffs from complex contracts). In the second case, an agent may fail to pay attention to a contractual offer. This seems a particularly pertinent problem in large organizations, where the decision to contract can only be taken once news of the offer has reached the right node of the organizational hierarchy. In either case, delays in readiness to trade should be expected especially in settings where gains from trading a particular good are difficult to anticipate in advance.³

Our model features a buyer and seller who each become ready to participate in a (potentially dynamic) contract at a random moment (we often refer to this moment as the "arrival date"). From the arrival date onwards, they remain ready and able to participate and communicate with the mechanism.⁴ Once both buyer and seller arrive, they can trade a single unit of a perishable good in each period. On arrival, the buyer and seller draw, respectively, a value for consuming the good and a cost of producing it, and these subsequently evolve over time according to a first-order Markov process. Efficient trade is trade which occurs if and only if both buyer and seller have arrived and the buyer's value exceeds the seller's cost.

Implementing efficient trade requires permitting buyer and seller participation in an efficient mechanism on any possible arrival date. Hence, unlike the aforementioned literature on repeated trade, an agent cannot be barred from participating only because he fails to do so at a particular instance. This makes participation constraints more difficult to satisfy, since an agent who just arrived has the option to wait and participate later, effectively mimicking later arrival. When values and costs evolve with time, this implies that even a buyer with the lowest value or a seller with the highest cost at his arrival date can still expect a positive rent in any mechanism implementing efficient trade.

We state simple necessary conditions for budget-balanced efficient trade (in an incentive-compatible mechanism satisfying all participation constraints). When the

³There are other reasons why agents may not be willing or able to participate, at least until after some delay. One is that regulatory compliance needs to be assured before parties can be confident that entering a trading relationship is legally permissible. Another is that the agents may need to be co-located in order to communicate or transact in any way.

⁴Random exogenous exits could, however, be easily accommodated.

buyer and seller arrive at each moment with positive probability, a necessary condition is that there be at least one date for which gains from trade are *certain* given that both the buyer and seller have arrived. In other words, the supports of the (marginal) distributions over values and costs at that date must not overlap. This may be seen as a dynamic analogue of the condition for a static environment due to Myerson and Satterthwaite (1983).

One way to understand this result is in view of the rents that the buyer and seller must be granted to ensure participation. As noted above, an agent who fails to participate necessarily retains the option to participate at a later date. Dissuading late participation then requires rents so large that (at least when the supports of values and costs always overlap and the traders may arrive at any moment) any efficient mechanism satisfying incentive-compatibility and participation constraints runs an expected budget deficit. The smallest feasible budget deficit turns out to equal that from a sequence of *static* mechanisms, each designed to ensure participation in the static mechanism alone. A possible interpretation is that, at least for reducing the expected budget deficit from efficient trade, truly dynamic mechanisms can be of limited use.⁵

The main outline of the paper is as follows. Following a discussion of the literature, Section 2 presents the model. Section 3 then presents our central results, while Section 4 provides a discussion of various caveats and extensions. Section 5 concludes. Proofs of all results are provided in the Appendix.

1.1 Related literature

As noted above, the key reference point for this paper is work on budget-balanced repeated trade when trading partners are known to be available to contract from the outset.⁶ This agenda was first developed in papers such as Athey and Miller (2007) and Athey and Segal (2007, 2013). While Athey and Miller consider values and costs drawn i.i.d. in each period, Athey and Segal consider persistent processes. Athey and Segal (2013) show, among other things, how to construct an efficient and budget-balanced "team mechanism" in which truthful strategies form a perfect Bayesian

⁵We show, however, that, when budget-balanced efficient trade is feasible, it may only be feasible through a dynamic mechanism in which participation at a given date gives rise to future obligations for the participants.

⁶The results in these papers continue to apply if agents instead arrive (stochastically) over time, but arrival dates are directly contractible.

equilibrium of a game in which agents' past reports are public information. When types follow an ergodic finite-state Markov process, they show that participation constraints can be satisfied provided players are sufficiently patient.

Subsequently, Skrzypacz and Toikka (2015) and Lamba (2013) provide further analysis of mechanisms implementing efficient trade while satisfying budget balance and individual rationality constraints. Notably, Skrzypacz and Toikka provide a necessary and sufficient condition for the existence of such mechanisms analogous to the condition developed for static problems by Makowski and Mezzetti (1994), Krishna and Perry (1998) and Williams (1999). This condition can be derived using payoff equivalence to a VCG mechanism: while a VCG mechanism for the bilateral trade problem runs an expected deficit equal to the expected gains from trade, each agent could be charged a participation fee to cover the deficit equal to the expected gains from trade conditional on their "worst" type. In a static mechanism, these fees do not cover the deficit unless type distributions do not overlap (an insight originally due to Myerson and Satterthwaite, 1983; see also Chatterjee and Samuelson, 1983, for a related analysis). In a dynamic setting with stochastically changing preferences (and commonly known arrival dates), however, participation fees can often cover the expected deficit if players are sufficiently patient. We make heavy use of such ideas in the analysis that follows, building especially on the observations of Skrzypacz and Toikka.

There has been a long held interest in allocation problems where players arrive over time. Recent examples include Board and Skrzypacz (2016) and Gershkov, Moldovanu and Strack (2016), who study "revenue management" problems, where revenue-maximizing mechanisms are designed that allocate goods to forward-looking buyers who arrive over time. While in these papers, buyers can perfectly anticipate their values from their arrival date onwards, Deb and Said (2015), Garrett (2016, forthcoming), and Ely, Garrett and Hinnosaar (forthcoming) consider settings in which buyers arrive over time, and these buyers also learn about their preferences over time. Our motivation in studying traders who arrive dynamically is closely related, but we focus on the question of implementing efficient allocations, rather than profit maximization.

Note that agents' private information on arrival times plays a central role in our setting, unlike what is seen in existing work on efficient dynamic mechanisms. For instance, Bergemann and Valimaki (2010) provide an efficient mechanism (the "dy-

dynamic pivot mechanism") which is ex-post incentive compatible, ex-post individually rational and satisfies efficient exit conditions. They note that their approach extends to the stochastic arrival of agents (privately informed of their arrival dates), assuming those who have not arrived are modeled as being in an "inactive state" in which their participation does not contribute to social surplus. Athey and Segal (2013, p 2477) similarly suggest that their approach extends readily to dynamic populations. The central message of our paper is different: agents having private information about arrival times can severely hamper the ability to implement efficient allocations. As explained above, this hinges on our view that an agent's "arrival" occurs on the first date the agent is able to contract rather than the first date at which an agent can contribute to surplus (as well as our desire for mechanisms that satisfy budget balance and induce agent participation in the mechanism at all arrival dates).⁷

More generally, our paper is related to the literature on mechanism design for agents whose preferences evolve stochastically (where the objective is often profit maximization rather than efficiency); see, among others, Baron and Besanko (1984), Besanko (1985), Courty and Li (2000), Battaglini (2005), Eso and Szentes (2007, forthcoming), Boleslavsky and Said (2013), Pavan, Segal and Toikka (2015), Battaglini and Lamba (2015) and Krahmer and Strausz (2015). Two connections to this literature are worth noting. First, we make direct use of insights in this literature, in particular by relying on the "dynamic payoff equivalence property", as formalized in Pavan, Segal and Toikka (see our Assumption 1 and the subsequent discussion below). Second, our result concerning the optimality of static mechanisms for reducing the budget deficit (while implementing efficient allocations) is reminiscent of Krahmer and Strausz's finding that static mechanisms are often profit-maximizing in settings with agent withdrawal rights. However, the reason for our result is different: it stems from agents' abilities to obtain efficient allocations even if they failed to contract in the past, rather than an ability to rescind the original contract.

⁷By interpreting an agent's "arrival" date as the first date at which he can contribute to surplus, we mean that each agent is available to contract at the beginning. Arrival would then be captured by supposing the buyer's value is initially too low for any trade, while the seller's cost is too high. The buyer would then "arrive" on the first date his value is high enough for efficient trade to be possible; for the seller, it would be the first date when his cost is sufficiently low. This case can be handled by the existing literature, see especially Proposition 1 of Skrzypacz and Toikka (2015). The same principles as in that literature suggest that efficient allocations could often be implemented with budget balance when agents are sufficiently patient: it would generally be enough that values and costs are not too persistent following each agent's random "arrival".

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 his arrival date, we suppose that the horizon is infinite.⁸ Periods are labeled $t = 1, 2, \dots$. Agents are labeled $i \in \{B, S\}$, with $-i$ denoting the potential trading partner of i . 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 and efficiency. In each period $t \geq \max\{\tau_B, \tau_S\}$, a period- t allocation $x_t \in \{0, 1\}$ is determined with $x_t = 1$ if the seller trades the good with the buyer. The resulting period- t payoff for the buyer is $\theta_{B,t}x_t + p_{B,t}$, where $\theta_{B,t}$ is the buyer's period- t value and $p_{B,t}$ is the date- t transfer paid to the buyer. The period- t payoff for the seller is $p_{S,t} - \theta_{S,t}x_t$, where $\theta_{S,t}$ is the seller's 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 $\theta_{B,t}$ and $\theta_{S,t}$ 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 for each agent i by $\theta_{i,s}^t = (\theta_{i,s}, \theta_{i,s+1}, \dots, \theta_{i,t})$.

Stochastic processes. Each agent $i \in \{B, S\}$ independently draws an arrival time τ_i from a distribution G_i with full support on the set of periods \mathbb{N} . Thus, let the probability that agent i arrives at date τ_i be $g_i(\tau_i) > 0$. 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 $\theta_{i,t} = \emptyset$ if agent i has not arrived by time t (although \emptyset is not a "payoff type").

The evolution of payoff types is also independent across agents. The set of possible (payoff) types at date t for agent i is denoted $\Theta_{i,t} = [\underline{\theta}_{i,t}, \bar{\theta}_{i,t}] \subset \mathbb{R}_+$. Assume the set $\cup_{t \geq 1} \Theta_{i,t}$ is bounded for each i . If an agent $i \in \{B, S\}$ arrives at any date τ_i , he draws at that date a type θ_{i,τ_i} from an absolutely continuous distribution $F_{\tau_i}^{i,In}(\theta_{i,\tau_i})$ which has full support on Θ_{i,τ_i} . Subsequently, at each date $t > \tau_i$, if the date $t - 1$ type is $\theta_{i,t-1} \in \Theta_{i,t-1}$, then he draws $\theta_{i,t}$ from an absolutely continuous conditional

⁸The results below are easily adapted to a finite horizon.

distribution $F_t^{i,Tr}(\theta_{i,t}|\theta_{i,t-1})$ with support on an interval $[\underline{\theta}_{i,t}(\theta_{i,t-1}), \bar{\theta}_{i,t}(\theta_{i,t-1})] \subset \Theta_{i,t}$ and a density $f_t^{i,Tr}(\theta_{i,t}|\theta_{i,t-1})$ strictly positive on this support. Assume further that each $F_t^{i,Tr}(\theta_{i,t}|\theta_{i,t-1})$ is continuous in $\theta_{i,t-1}$ uniformly across $(\theta_{i,t-1}, \theta_{i,t}) \in \Theta_{i,t-1} \times \Theta_{i,t}$ for each i, t .

The above description of the process encodes our assumption that (payoff) types evolve according to a (possibly time-varying) first-order Markov process. The role of the restriction that the support of $F_t^{i,Tr}(\theta_{i,t}|\theta_{i,t-1})$ be contained in $\Theta_{i,t}$ will be discussed in detail in Section 4.2. For now, note that it is arguably quite mild, since it is implied if we take $F_{\tau_i}^{i,In}$, for each i and $\tau_i \geq 2$, to be the marginal distribution at date τ_i of the payoff-type process conditional on arrival at date 1. In this case, we can view each agent i 's payoff types as following a common (latent) process from date 1, with the arrival time τ_i determined independently of this 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, Segal and Toikka, 2014, as well as Skrzypacz and Toikka), and we give a sufficient condition below. Many stochastic processes will satisfy our conditions. A commonly used example is the first-order autoregressive process $\theta_{i,t} = \gamma\theta_{i,t-1} + (1 - \gamma)m_i + \varepsilon_{i,t}$, with $\gamma \in (0, 1)$, $m_i \in \mathbb{R}_+$ the long-run mean of $\theta_{i,t}$, and $\varepsilon_{i,t}$ a mean-zero random variable, with appropriate regularity conditions on the distribution of the initial type and subsequent innovations $\varepsilon_{i,t}$.

Mechanisms. Without loss of generality, we study direct mechanisms. Each agent i makes a report of his (payoff) type $\hat{\theta}_{i,\hat{\tau}_i} \in \Theta_{i,\hat{\tau}_i}$ on the first date of participation $\hat{\tau}_i$, and then continues to provide updates of these types at each date. In particular, if agent i reported $\hat{\theta}_{i,t-1}$ at date $t - 1$, then he is permitted a report in the support of $F_t^{i,Tr}(\cdot|\hat{\theta}_{i,t-1})$ at date t . The reports of each agent i up to date t may then be denoted $\hat{\theta}_{i,\hat{\tau}_i}^t$. Here, given that $\hat{\tau}_i$ is the first date at which agent i reports to the mechanism, it may be thought of as i 's reported arrival time.

A direct mechanism $\Omega = \langle x_t, p_{B,t}, p_{S,t} \rangle_{t \geq 1}$ 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 $x_t(\hat{\theta}_{B,\hat{\tau}_B}^t, \hat{\theta}_{S,\hat{\tau}_S}^t) \in \{0, 1\}$ for each possible pair of report se-

quences $(\hat{\theta}_{B,\hat{\tau}_B}^t, \hat{\theta}_{S,\hat{\tau}_S}^t) \in \Pi_{s=\hat{\tau}_B}^t \Theta_{B,s} \times \Pi_{s=\hat{\tau}_S}^t \Theta_{S,s}$.⁹ Date- t payments to agent i are $p_{i,t}(\hat{\theta}_{i,\hat{\tau}_i}^t, \hat{\theta}_{-i,\hat{\tau}_{-i}}^t) \in \mathbb{R}$ for reports $(\hat{\theta}_{i,\hat{\tau}_i}^t, \hat{\theta}_{-i,\hat{\tau}_{-i}}^t)$, where $-i$ indicates the agent other than i . Both allocations and transfers are assumed throughout to be measurable functions of the agents' reports. Note here that the length of the report sequences that are arguments to the payment and allocation rules indicate the arrival times of the agents (for instance, if $\hat{\tau}_B = \hat{\tau}_S = t$ for some t , then $p_{i,t}(\hat{\theta}_{i,\hat{\tau}_i}^t, \hat{\theta}_{-i,\hat{\tau}_{-i}}^t) = p_{i,t}(\hat{\theta}_{i,t}, \hat{\theta}_{-i,t})$ give the first payments received by each agent i , since the length of each report sequence is 1).

It may now be helpful to delineate the timing of events in each period. At the beginning of each period t , each agent i has made a sequence of reports $\hat{\theta}_{i,\hat{\tau}_i}^{t-1} \in \Pi_{s=\hat{\tau}_i}^{t-1} \Theta_{i,s}$ if the agent already participated in the mechanism at some date $\hat{\tau}_i < t$. If agent i arrived before date t , then he (privately) draws a date- t payoff type $\theta_{i,t}$ from the distribution $F_t^{i,Tr}(\theta_{i,t}|\theta_{i,t-1})$, where $\theta_{i,t-1}$ denotes agent i 's true (payoff) type at $t-1$. If he arrives at date t , then he (privately) draws $\theta_{i,t}$ from $F_t^{i,In}$. He then (simultaneously with agent $-i$) makes a report $\hat{\theta}_{i,t} \in \Theta_{i,t}$ to the mechanism. If this is the first report, then it amounts to a claim that the arrival time is $\hat{\tau}_i = t$ (and $\hat{\tau}_i$ is then agent i 's reported arrival time). At this point, the allocation $x_t(\hat{\theta}_{B,\hat{\tau}_B}^t, \hat{\theta}_{S,\hat{\tau}_S}^t)$ can be determined and the transfers $p_{i,t}(\hat{\theta}_{i,\hat{\tau}_i}^t, \hat{\theta}_{-i,\hat{\tau}_{-i}}^t)$ 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 $\hat{\theta}_{-i,\hat{\tau}_{-i}}^{t-1}$. As Myerson (1986) noted, incentive constraints are most easily satisfied when agents are *least* informed. However, realistically in most settings of interest, some information "leaks" through agents observing outcomes, their payoffs, or both. This creates the potential complication of (i) specifying precisely what agents observe (the outcome x_t at each t , the transfers of different agents, and/or their reports), and (ii) deducing mechanisms which most effectively hide information from agents while implementing

⁹Random allocations $x_t(\hat{\theta}_{B,\hat{\tau}_B}^t, \hat{\theta}_{S,\hat{\tau}_S}^t) \in (0,1)$ could readily be permitted, but are not needed in what follows.

the efficient allocation in an incentive-compatible manner.

Avoiding these complications, we focus on so-called "*blind*" mechanisms in which each agent never receives any information about the other's past reports. This corresponds to an environment in which neither trade outcomes nor payoffs are observed. Our focus makes sense in light of Myerson's (1986) observation and because we are chiefly interested in negative results (for instance, we will use that, if budget-balanced efficient trade cannot be supported in the blind mechanism, then the same is true in any mechanism with information "leakage"). Nonetheless, where blind mechanisms exist satisfying various desiderata (efficiency, incentive compatibility, willingness to participate, budget neutrality, etc.), it is of interest to understand whether we can also satisfy the same desiderata when each agent is informed about the other's past decisions. Of particular interest is the "*public mechanism*" in which all past reports are publicly revealed (and available to any agent, irrespective of his arrival date). If the public mechanism satisfies the relevant incentive constraints, then any less informative mechanism will do so as well (with a less informative mechanism, the relevant constraints are merely "pooled" and hence continue to be satisfied).

Agent continuation payoffs. Consider a public mechanism Ω . The buyer's expected continuation payoff in Ω when reporting truthfully after a history of reports $(\theta_{B,\tau_B}^{t-1}, \theta_{S,\tau_S}^{t-1})$ (if any), when his date- t value is $\theta_{B,t}$, is¹⁰

$$\begin{aligned} & V_{B,t}^\Omega(\theta_{B,\tau_B}^t; \theta_{S,\tau_S}^{t-1}) \\ = & \mathbb{E} \left[\sum_{s=t}^{\infty} \delta^{s-t} \begin{pmatrix} p_{B,s}(\tilde{\theta}_{B,\tau_B}^s, \tilde{\theta}_{S,\tau_S}^s) \\ + \tilde{\theta}_{B,s} x_s(\tilde{\theta}_{B,\tau_B}^s, \tilde{\theta}_{S,\tau_S}^s) \end{pmatrix} \middle| \begin{matrix} \tilde{\theta}_{B,\tau_B}^t = \theta_{B,\tau_B}^t \\ \tilde{\theta}_{S,\tau_S}^{t-1} = \theta_{S,\tau_S}^{t-1} \end{matrix} \right]. \end{aligned} \quad (1)$$

Analogously, if the seller's date- t cost is $\theta_{S,t}$, then his expected continuation value is

$$\begin{aligned} & V_{S,t}^\Omega(\theta_{S,\tau_S}^t; \theta_{B,\tau_B}^{t-1}) \\ = & \mathbb{E} \left[\sum_{s=t}^{\infty} \delta^{s-t} \begin{pmatrix} p_{S,s}(\tilde{\theta}_{S,\tau_S}^s, \tilde{\theta}_{B,\tau_B}^s) \\ - \tilde{\theta}_{S,s} x_s(\tilde{\theta}_{B,\tau_B}^s, \tilde{\theta}_{S,\tau_S}^s) \end{pmatrix} \middle| \begin{matrix} \tilde{\theta}_{S,\tau_S}^t = \theta_{S,\tau_S}^t \\ \tilde{\theta}_{B,\tau_B}^{t-1} = \theta_{B,\tau_B}^{t-1} \end{matrix} \right]. \end{aligned} \quad (2)$$

In the blind mechanism, the other agent's reports are not available, and so we simply write $V_{B,t}^\Omega(\theta_{B,\tau_B}^t)$ and $V_{S,t}^\Omega(\theta_{S,\tau_S}^t)$. By the law of iterated expectations, we have

¹⁰We use tildes to denote random variables.

$$V_{i,t}^{\Omega}(\theta_{i,\tau_i}^t) = \mathbb{E} \left[V_{i,t}^{\Omega} \left(\theta_{i,\tau_i}^t; \tilde{\theta}_{-i,\tau-i}^{t-1} \right) \right] \text{ for each agent } i.$$

Incentive compatibility. Our notion of incentive compatibility depends on whether the mechanism is blind or public. To define incentive compatibility, suppose that each agent i reports to the mechanism on the arrival date τ_i . A blind mechanism is then *Bayesian incentive compatible (BIC)* if, for each i and participation date τ_i , his expected payoffs are maximized by reporting payoff types truthfully (thus $V_{i,\tau_i}^{\Omega}(\theta_{i,\tau_i})$ is equal to the supremum of his expected continuation payoff over all possible reporting strategies, for each i , each τ_i , and each $\theta_{i,\tau_i} \in \Theta_{i,\tau_i}$). A public mechanism is *perfect-Bayesian incentive compatible (PIC)* if the reporting game specified above, assuming both agents participate at their arrival dates,¹¹ has a perfect-Bayesian equilibrium in which each agent reports truthfully upon arrival and then continues to report truthfully provided he was truthful in the past. Note here that, having restricted the space of reports (namely, to the support of the conditional distribution $F_t^{i,Tr}(\cdot|\hat{\theta}_{i,t-1})$ for each i , date t , and previous report $\hat{\theta}_{i,t-1} \in \Theta_{i,t-1}$), no report sequence is ever "off-path".¹² Hence, each agent's beliefs over the other agent's types in the public mechanism are simply those given by truthful reporting with probability one.

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. Clearly, this only strengthens our results regarding the absence of satisfactory mechanisms for efficient trade.

Note that each agent i 's payoff type $\theta_{i,t}$ is a sufficient statistic for the evolution of i 's future types. Hence, and given our restriction on the supports of $(F_t^{i,Tr})_{t \geq 2}$, an agent who delays participation (say by one period) finds himself in the same situation as if he arrived at a later date (the date after his true arrival date).

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). By the one-shot deviation principle, it is enough to check deviations in which each agent delays participation by one period. In the

¹¹Hence, both our notions of incentive compatibility relate only to reports of payoff types and not arrival times at the mechanism. Constraints on agents' willingness to participate in (incentive-compatible) mechanisms are considered next.

¹²See Skrzypacz and Toikka (2015) for the same assumption.

public mechanism, such a deviation is never profitable for agent i at date τ_i if

$$V_{i,\tau_i}^\Omega \left(\theta_{i,\tau_i}; \theta_{-i,\tau_i}^{\tau_i-1} \right) \geq \delta \mathbb{E} \left[V_{i,\tau_i+1}^\Omega \left(\tilde{\theta}_{i,\tau_i+1}; \tilde{\theta}_{-i,\tilde{\tau}_{-i}}^{\tau_i} \right) \mid \tilde{\theta}_{i,\tau_i} = \theta_{i,\tau_i}, \tilde{\theta}_{-i,\tilde{\tau}_{-i}}^{\tau_i-1} = \theta_{-i,\tau_i}^{\tau_i-1} \right] \quad (3)$$

for all $\theta_{i,\tau_i} \in \Theta_{i,\tau_i}$ and all past reports $\theta_{-i,\tau_i}^{\tau_i-1}$ for agent $-i$. In the blind mechanism, the deviation is never profitable provided simply that

$$V_{i,\tau_i}^\Omega \left(\theta_{i,\tau_i} \right) \geq \delta \mathbb{E} \left[V_{i,\tau_i+1}^\Omega \left(\tilde{\theta}_{i,\tau_i+1} \right) \mid \tilde{\theta}_{i,\tau_i} = \theta_{i,\tau_i} \right] \quad (4)$$

for all $\theta_{i,\tau_i} \in \Theta_{i,\tau_i}$. We say that a mechanism in which (3) always holds satisfies "public participation constraints" (or *PPC*), while if (4) always holds it satisfies "blind participation constraints" (or *BPC*).

Budget balance. A mechanism is said to be *budget balanced* (*BB*) if $p_{B,t} \left(\theta_{B,\tau_B}^t, \theta_{S,\tau_S}^t \right) + p_{S,t} \left(\theta_{S,\tau_S}^t, \theta_{B,\tau_B}^t \right) = 0$ for all report sequences $\left(\theta_{B,\tau_B}^t, \theta_{S,\tau_S}^t \right)$. It runs an *expected budget surplus* (or satisfies *EBS*) if

$$U_1 \equiv -\mathbb{E} \left[\sum_{t=1}^{\infty} \delta^{t-1} p_{B,t} \left(\tilde{\theta}_{B,\tilde{\tau}_B}^t, \tilde{\theta}_{S,\tilde{\tau}_S}^t \right) + \sum_{t=1}^{\infty} \delta^{t-1} p_{S,t} \left(\tilde{\theta}_{S,\tilde{\tau}_S}^t, \tilde{\theta}_{B,\tilde{\tau}_B}^t \right) \right] \quad (5)$$

is non-negative.

Here, U_1 is the ex-ante profit of a "third-party broker", in the spirit of the broker introduced by Myerson and Satterthwaite (1983), so we refer to U_1 as the "broker's expected surplus". Recall that the strongest version of Myerson and Satterthwaite's impossibility result for the static environment was that the third-party broker cannot break even in a mechanism implementing efficient trade. Likewise, one of our main concerns is whether efficient trade can be sustained while satisfying EBS.

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

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 defined by $F_{\tau_i}^{i,In}$ and $F_t^{i,Tr}$ for all $i \in \{B, S\}$, and all periods τ_i and $t \geq 2$, satisfy the "payoff-equivalence property" meaning that*

the following holds. Consider any two BIC blind mechanisms $\Omega = \langle x_t, p_{B,t}, p_{S,t} \rangle_{t \geq 1}$ and $\Omega' = \langle x'_t, p'_{B,t}, p'_{S,t} \rangle_{t \geq 1}$ satisfying $x_t = x'_t$ for all t . There exist real-valued scalars $(b_{i,\tau_i})_{i \in \{B,S\}, \tau_i \geq 1}$ such that, for each agent i , and each date τ_i , $V_{i,\tau_i}^\Omega(\theta_{i,\tau_i}) = b_{\tau_i} + V_{i,\tau_i}^{\Omega'}(\theta_{i,\tau_i})$ for all $\theta_{i,\tau_i} \in \Theta_{i,\tau_i}$.

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. This need not guarantee the same is true for the public mechanism, where the other agent's reports up to τ_i are revealed, even permitting that the constants b_{i,τ_i} can depend on the information revealed prior to τ_i . While we could state an analogous payoff equivalence property for the public environment, this turns out not to be necessary for what follows.

While more general conditions are available (see, especially, Pavan, Segal and Toikka, 2014), payoff equivalence holds if the following two criteria are satisfied: (i) there exists, for each $i \in \{B, S\}$, a sequence of continuously differentiable functions $(z_{i,t})_{t \geq 2}$, with $z_{i,t} : \Theta_{i,t-1} \times [0, 1] \rightarrow \Theta_{i,t}$ such that, when $\tilde{\varepsilon}$ is uniformly distributed on $[0, 1]$, $z_{i,t}(\theta_{i,t-1}, \tilde{\varepsilon})$ is distributed according to the conditional distribution $F_t^{i,Tr}(\cdot | \theta_{i,t-1})$, and (ii) there exists $k < \frac{1}{\delta}$ such that $\left| \frac{\partial z_{i,t}(\theta_{i,t-1}, \varepsilon)}{\partial \theta_{i,t-1}} \right| \leq k$ for all $(\theta_{i,t-1}, \varepsilon) \in \Theta_{i,t-1} \times [0, 1]$. The condition implies that the distribution of an agent's date- t payoff type is not too sensitive to the realization of his previous period's payoff type.

No Ponzi schemes. While the above requirements are natural analogues of conditions considered in the earlier literature, one additional restriction on mechanisms (already implicit in the above) is worth emphasizing. In particular, we restrict attention throughout to mechanisms such that the expression (5) is well-defined, and refer to this condition as "*no Ponzi schemes*" (or NPS). This terminology is justified on the grounds it implies that the expected discounted payoffs of the broker, the buyer and the seller are well-defined and sum to the expected discounted surplus from trade. For an efficient blind mechanism Ω , this is the statement that

$$\begin{aligned} & U_1 + \mathbb{E} \left[\delta^{\tilde{\tau}_B - 1} V_{B, \tilde{\tau}_B}^\Omega \left(\tilde{\theta}_{B, \tilde{\tau}_B} \right) \right] + \mathbb{E} \left[\delta^{\tilde{\tau}_S - 1} V_{S, \tilde{\tau}_S}^\Omega \left(\tilde{\theta}_{S, \tilde{\tau}_S} \right) \right] \\ &= \mathbb{E} \left[\sum_{t=\max\{\tilde{\tau}_B, \tilde{\tau}_S\}}^{\infty} \delta^{t-1} x^E \left(\tilde{\theta}_{B,t}, \tilde{\theta}_{S,t} \right) \left(\tilde{\theta}_{B,t} - \tilde{\theta}_{S,t} \right) \right]. \end{aligned} \quad (6)$$

Condition NPS will be important for, in the blind environment, there exist BB mechanisms failing NPS which guarantee arbitrary values of the expected payoffs from participation, $V_{i,\tau_i}^\Omega(\theta_{i,\tau_i})$, essentially permitting participation constraints to be satisfied "for free". Given that the mechanism is BB, no third-party broker is required to facilitate trade, so it arguably makes sense that the calculation in (5) is redundant.

To see how a Ponzi scheme can work in the absence of NPS, we could start with a mechanism satisfying E, BB and BIC, but not necessarily BPC (suppose any participant faces a sequence of AGV mechanisms, following d'Aspremont and Gérard-Varet, 1979). Then note that any level of the expected continuation payoff $V_{i,\tau_i}^\Omega(\theta_{i,\tau_i})$ could be assured for any agent i participating at τ_i through an appropriate system of type-independent transfers, while leaving payoffs for agent i unchanged for other participation dates, and the other agent $-i$'s payoffs unchanged at all dates. In particular, suppose we want to increase $V_{i,\tau_i}^\Omega(\theta_{i,\tau_i})$ by one unit, holding the payoffs $V_{i,\hat{\tau}_i}^\Omega(\theta_{i,\hat{\tau}_i})$ constant for all other participation dates $\hat{\tau}_i \neq \tau_i$, and holding the other agent's payoffs $V_{-i,\tau_{-i}}^\Omega(\theta_{-i,\tau_{-i}})$ constant for all participation dates τ_{-i} . Then, we could require agent $-i$ to pay agent i the amount $\frac{1}{\delta g_{-i}(\tau_i+1)}$ if agent $-i$ arrives at $\tau_i + 1$. If $-i$ arrives at date $\tau_i + 1$, then his expected additional payment to i is $\frac{g_i(\tau_i)}{\delta g_{-i}(\tau_i+1)}$. Thus, if $-i$ arrives at date $\tau_i + 1$ and i at date $\tau_i + 2$, require i to pay to $-i$ the amount $\frac{g_i(\tau_i)}{\delta g_{-i}(\tau_i+1)}$ (hence, the expected payoff of $-i$ remains unchanged relative to the original mechanism). We can then ask agent $-i$ to compensate agent i in case i arrives at $\tau_i + 2$ and $-i$ at $\tau_i + 3$, and so forth.

There are several reasons to impose NPS. First, for mechanisms satisfying BB but failing NPS, agents' ex-ante expected payments are not well-defined, and this might be viewed as rendering agents' ex-ante expected payoffs ambiguous. This could be important, for instance, if agents are required to take some ex-ante decision (say at date zero) as to whether to remain "receptive" to trade (for instance, agents may need to incur some initial cost in order to arrive to the mechanism at the rates implied by G_i , $i \in \{B, S\}$). Second, there is considerable precedent in the literature for considering condition EBS (recall, for instance, Myerson and Satterthwaite's, 1983, concern with expected broker surplus, and Athey and Miller's, 2007, concern with actuarially fair insurance). However, the condition EBS only makes sense if U_1 in (5) is well-defined. Third, if NPS fails, then the expected magnitude of transfers at each date cannot be uniformly bounded, and this may seem an undesirable feature of any

trading mechanism.¹³ Fourth, the Ponzi schemes envisaged above do not work in the environment where agents' reports are public (and we ask for mechanisms satisfying BB, E, PIC and PPC).¹⁴ This follows because, once an agent i is commonly known to have participated at τ_i , any payments to him from $-i$ at dates $\tau_i + 1$ or later must either affect the expected present value of $-i$'s payoffs or be compensated by i conditional on having arrived at τ_i (hence leaving his expected payoff conditional on arrival at τ_i unchanged). Hence, our impossibility result will apply to efficient trade in the public environment subject to BB, PIC and PPC without additional restrictions on transfers.¹⁵

3.2 Main results

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

Lemma 1 *If a blind mechanism Ω maximizes the broker's expected surplus U_1 among mechanisms satisfying E, BIC and BPC, the following condition holds for each $i \in \{B, S\}$ and each date $\tau_i \in \mathbb{N}$:*

$$\inf_{\theta_{i,\tau_i} \in \Theta_{i,\tau_i}} \left\{ V_{i,\tau_i}^\Omega(\theta_{i,\tau_i}) - \delta \mathbb{E} \left[V_{i,\tau_i+1}^\Omega(\tilde{\theta}_{i,\tau_i+1}) \mid \tilde{\theta}_{i,\tau_i} = \theta_{i,\tau_i} \right] \right\} = 0.$$

The reason for this result is simple. At date τ_i , among mechanisms satisfying BIC and inducing agent i 's participation at date $\tau_i + 1$, agent i 's date- τ_i participation constraint is given by (4). This states that agent i must prefer to participate at date τ_i than delay until $\tau_i + 1$, taking his chances to participate with the new realized payoff type. Analogously to static mechanism design, minimizing agent expected rents then requires this participation constraint to bind. Indeed, this must be the case in a mechanism that maximizes U_1 , since lowering $V_{i,\tau_i}^\Omega(\theta_{i,\tau_i})$ does not interfere with willingness to participate at dates $\tau_i - 1$ or earlier.

We now show that, to maximize the broker's surplus, it is enough to rely on a sequence of static mechanisms. We begin by considering the static VCG mechanism which ensures each agent a payoff equal to the surplus from trade, i.e.

¹³The mechanisms we construct for our positive results below satisfy NPS and do not have this feature.

¹⁴More formally, the analogue of Condition (6) for the public environment must hold in any incentive-compatible mechanism (where agent payoffs at participation, as defined by (1) and (2), are always well-defined).

¹⁵The same is true for our weaker budgetary condition, EBS.

$(\theta_{B,t} - \theta_{S,t}) x^E(\theta_{B,t}, \theta_{S,t})$ for date t . At date t , the lowest realization of the expected surplus for an agent i over date- t payoff types $\theta_{i,t} \in \Theta_{i,t}$, given only that the other agent $-i$ has arrived by date t , is

$$\inf_{\theta_{i,t} \in \Theta_{i,t}} \mathbb{E} \left[\left(\tilde{\theta}_{B,t} - \tilde{\theta}_{S,t} \right) x^E \left(\tilde{\theta}_{B,t}, \tilde{\theta}_{S,t} \right) \mid \tilde{\theta}_{i,t} = \theta_{i,t}, \tilde{\theta}_{-i,t} \neq \emptyset \right]. \quad (7)$$

The lowest realizations occur for the lowest value of the buyer and highest cost of the seller. Assuming payments occur only when both agents participate in the mechanism, then agents could be charged this minimal surplus as a fixed participation fee. In particular, they would still be willing to participate and report truthfully in the static (date- t) mechanism for all realizations of $\theta_{i,t}$ (here, assuming that participation in the date- t static mechanism does not affect the opportunities or obligations of the agents at future dates). We then define (perhaps abusively) a "VCG-* 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

$$\begin{aligned} p_{B,t}^{VCG-*}(\theta_{B,\tau_B}^t, \theta_{S,\tau_S}^t) &= -\theta_{S,t} x^E(\theta_{B,t}, \theta_{S,t}) \\ &\quad - \mathbb{E} \left[\left(\underline{\theta}_{B,t} - \tilde{\theta}_{S,t} \right) x^E \left(\underline{\theta}_{B,t}, \tilde{\theta}_{S,t} \right) \mid \tilde{\theta}_{S,t} \neq \emptyset \right] \end{aligned} \quad (8)$$

for the buyer, and

$$\begin{aligned} p_{S,t}^{VCG-*}(\theta_{S,\tau_S}^t, \theta_{B,\tau_B}^t) &= \theta_{B,t} x^E(\theta_{B,t}, \theta_{S,t}) \\ &\quad - \mathbb{E} \left[\left(\tilde{\theta}_{B,t} - \bar{\theta}_{S,t} \right) x^E \left(\tilde{\theta}_{B,t}, \bar{\theta}_{S,t} \right) \mid \tilde{\theta}_{B,t} \neq \emptyset \right] \end{aligned} \quad (9)$$

for the seller. If one or both agents have not arrived by date t , then there is no trade and no transfers are made. We then find the following.

Proposition 1 *The broker's expected surplus U_1 is maximized (among mechanisms satisfying E, BIC and BPC) by running a blind VCG-* mechanism at each date t ; that is, by setting $\Omega = (x^E, p_{B,t}^{VCG-*}, p_{S,t}^{VCG-*})_{t \geq 1}$. The same value U_1 is attainable in a public mechanism that is E, PIC and PPC.*

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. If the transfers are augmented by fixed (i.e., type-independent) participation fees as in (8) and (9),

then, at each date t , the buyer with the lowest value $\underline{\theta}_{B,t}$ and the seller with the highest cost $\bar{\theta}_{S,t}$ obtain an additional expected payoff of zero by participating in the date- t mechanism, rather than delaying participation to the subsequent period. Hence participation constraints always bind. Using the "payoff equivalence property" of Assumption 1, we then show that the sequence of VCG-* mechanisms generates the highest expected surplus for the broker among mechanisms implementing efficiency and satisfying the condition of Lemma 1 at each date t . The result for public mechanisms is then shown by describing a public mechanism (also a sequence of static VCG mechanisms, but with different fees) satisfying E, PIC and PPC, and which attains the same value of U_1 as the aforementioned blind mechanism.¹⁶

What is the broker's surplus U_1 in the broker-optimal efficient mechanism? First, recall that the VCG mechanism that we took as our starting point runs a budget deficit in the bilateral trade problem equal to the total surplus $(\theta_{B,t} - \theta_{S,t}) x^E(\theta_{B,t}, \theta_{S,t})$ at each date t . However, the VCG-* mechanisms collect fees equal to (7) from each agent i . We therefore have the following result.

Proposition 2 *The largest value of the broker's expected surplus in a blind mechanism satisfying E, BIC and BPC (alternatively, in a public mechanism satisfying E, PIC and PPC) is*

$$U_1 = \sum_{t=1}^{\infty} \delta^{t-1} G_B(t) G_S(t) \Psi_t, \quad (10)$$

where

$$\Psi_t = \begin{aligned} & -\mathbb{E} \left[\left(\tilde{\theta}_{B,t} - \tilde{\theta}_{S,t} \right) x^E \left(\tilde{\theta}_{B,t}, \tilde{\theta}_{S,t} \right) \mid \tilde{\theta}_{B,t} \neq \emptyset, \tilde{\theta}_{S,t} \neq \emptyset \right] \\ & + \mathbb{E} \left[\left(\underline{\theta}_{B,t} - \tilde{\theta}_{S,t} \right) x^E \left(\underline{\theta}_{B,t}, \tilde{\theta}_{S,t} \right) \mid \tilde{\theta}_{S,t} \neq \emptyset \right] \\ & + \mathbb{E} \left[\left(\tilde{\theta}_{B,t} - \bar{\theta}_{S,t} \right) x^E \left(\tilde{\theta}_{B,t}, \bar{\theta}_{S,t} \right) \mid \tilde{\theta}_{B,t} \neq \emptyset \right] \end{aligned} . \quad (11)$$

There exists a blind mechanism that satisfies BB, E, BIC and BPC if and only if $U_1 \geq 0$. A sufficient condition for $U_1 < 0$ is that $\Theta_{B,t} \cap \Theta_{S,t}$ has positive length for each t , meaning that the distributions $\Pr \left(\tilde{\theta}_{B,t} \leq \theta_{B,t} \mid \tilde{\theta}_{B,t} \neq \emptyset \right)$ and $\Pr \left(\tilde{\theta}_{S,t} \leq \theta_{S,t} \mid \tilde{\theta}_{S,t} \neq \emptyset \right)$ have "overlapping supports" for all t .

Equation (10) should be understood as a weighted average of the broker's maximal expected surplus under efficient trade in static mechanisms, as calculated, for instance,

¹⁶That this is enough follows because the broker's surplus can be no higher in a public mechanism satisfying PIC and PPC than in a blind mechanism with the same allocation rule satisfying BIC and BPC.

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 . Indeed, this corresponds to the broker's expected surplus in a VCG-* mechanism as defined above. The weights in (10) comprise the discount factor δ^{t-1} and the probability that both agents arrive by date t , $G_B(t)G_S(t)$.¹⁷

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 (whether in blind or public mechanisms). Conversely, if $\Psi_t \geq 0$ at a given date t , then an efficient and budget-balanced static mechanism exists at date t ; indeed, gains from trade must be assured conditional on both agents arriving by date t , and so efficient trade 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 static (posted-price) mechanisms.

If $\Psi_t < 0$ for some t and yet the expression in (10) is non-negative, then a blind mechanism can be chosen to satisfy E, BB, BIC and BPC. The mechanism constructed 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 , some of these additional payments must 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 periods, 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 U_1 , truly "dynamic" mechanisms are needed to obtain budget balance.)

Finally, note that Proposition 2 does not extend immediately to the existence of public mechanisms satisfying BB, E, PIC and PPC. Consider the case where $U_1 \geq 0$, and where $\Psi_1 > 0$ but $\Psi_t < 0$ for all $t \geq 2$. If the information that no agent arrived at date 1 becomes public, then the continuation mechanism from date-2 onwards must

¹⁷Note that the broker's expected surplus (10) thus depends only on the distributions of arrival times and the marginal distributions of payoff types, $\Pr(\tilde{\theta}_{B,t} \leq \theta_{B,t} | \tilde{\theta}_{B,t} \neq \emptyset)$ and $\Pr(\tilde{\theta}_{S,t} \leq \theta_{S,t} | \tilde{\theta}_{S,t} \neq \emptyset)$, 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.

be run with an expected loss for the broker (assuming it satisfies E, PIC and PPC), and so budget-balanced efficient trade is not feasible. This observation stands in contrast to the environment where agents arrive at the outset with probability one, and so there is a single constraint for each agent's participation at the initial date. Then, Skrzypacz and Toikka (2015; Proposition 1, cases (ii) and (iii)) show that, if there exists a blind mechanism that satisfies E, EBS, BIC and the initial participation constraint for each agent, then there exists a public mechanism satisfying E, BB, PIC and the same participation constraints.¹⁸

4 Extensions and discussion

4.1 One-sided uncertainty

The above insights can be readily adapted to settings in which arrival is known not to occur at some dates. For instance, consider the case where the seller is commonly known to be in the market at date 1 (similar results hold when instead the buyer is commonly known to be present at date 1). Then we have the following analogue of Proposition 2.

Proposition 3 *Suppose that the seller is commonly known to be in the market at date 1, but that the buyer may arrive at any date (i.e., $G_B(\cdot)$ continues to have full support on \mathbb{N}). The largest value of the broker's expected surplus in a blind mechanism satisfying E, BIC and BPC is*

$$\bar{U}_1 = \sum_{t=1}^{\infty} \delta^{t-1} G_B(t) \bar{\Psi}_t, \quad (12)$$

where

$$\begin{aligned} \bar{\Psi}_t = & -\mathbb{E} \left[\left(\tilde{\theta}_{B,t} - \tilde{\theta}_{S,t} \right) x^E \left(\tilde{\theta}_{B,t}, \tilde{\theta}_{S,t} \right) \mid \tilde{\theta}_{B,t} \neq \emptyset \right] \\ & + \mathbb{E} \left[\left(\underline{\theta}_{B,t} - \tilde{\theta}_{S,t} \right) x^E \left(\underline{\theta}_{B,t}, \tilde{\theta}_{S,t} \right) \right] \\ & + \mathbb{E} \left[\left(\tilde{\theta}_{B,t} - \tilde{\theta}_{S,t} \right) x^E \left(\tilde{\theta}_{B,t}, \tilde{\theta}_{S,t} \right) \mid \tilde{\theta}_{B,t} \neq \emptyset, \tilde{\theta}_{S,1} = \bar{\theta}_{S,1} \right] \end{aligned} \quad (13)$$

A mechanism satisfying BB, E, BIC and BPC exists if and only if $\bar{U}_1 \geq 0$. In case $\bar{U}_1 \geq 0$, there exists also a public mechanism satisfying BB, E, PIC and PPC.

¹⁸The implication in Skrzypacz and Toikka does extend to our setting, however, if just one of the agents has a commonly known arrival date, which is the case we consider next.

The expression (12) can be understood as follows. Consider first running a (static) VCG mechanism at each date, such that each agent earns a payoff equal to the surplus from trade. Both agents will participate in these mechanisms whenever possible and report truthfully, so the ex-ante expected present value of the budget deficit is

$$\sum_{t=1}^{\infty} \delta^{t-1} G_B(t) \mathbb{E} \left[\left(\tilde{\theta}_{B,t} - \tilde{\theta}_{S,t} \right) x^E \left(\tilde{\theta}_{B,t}, \tilde{\theta}_{S,t} \right) \mid \tilde{\theta}_{B,t} \neq \emptyset \right].$$

Now suppose that the buyer is charged a fixed participation fee equal to $\mathbb{E} \left[\left(\underline{\theta}_{B,t} - \tilde{\theta}_{S,t} \right) x^E \left(\underline{\theta}_{B,t}, \tilde{\theta}_{S,t} \right) \right]$ for participation in the VCG mechanism at each date. At each date t , the buyer is willing to participate (given that the mechanism is blind and participation does not affect the mechanism he faces at any future date), and he is indifferent for the lowest value realization $\underline{\theta}_{B,t}$. If the seller has the highest date-1 cost $\bar{\theta}_{S,1}$, then his expected payoff in the sequence of VCG mechanisms is

$$\sum_{t=1}^{\infty} \delta^{t-1} G_B(t) \mathbb{E} \left[\left(\tilde{\theta}_{B,t} - \tilde{\theta}_{S,t} \right) x^E \left(\tilde{\theta}_{B,t}, \tilde{\theta}_{S,t} \right) \mid \tilde{\theta}_{B,t} \neq \emptyset, \tilde{\theta}_{S,1} = \bar{\theta}_{S,1} \right]. \quad (14)$$

This type of the seller is just willing to participate if charged a participation fee equal to (14), given that he is permitted to participate in the mechanism only if participating at date 1 (and hence his payoff if declining to participate at date 1 is zero). All lower cost types are then willing to participate as well. Expression (12) then comprises the budget deficit of the original VCG mechanisms, plus the sum of the expected fees.

Compared to the case with uncertain arrival for both agents, efficient trade is easier to sustain while satisfying EBS, since there are fewer participation constraints. For instance, EBS may be satisfied (say, in a blind mechanism that is E, BIC and BPC) even if the supports of $\Pr \left(\tilde{\theta}_{B,t} \leq \theta_{B,t} \mid \tilde{\theta}_{B,t} \neq \emptyset \right)$ and $\Pr \left(\tilde{\theta}_{S,t} \leq \theta_{S,t} \right)$ overlap at each date t . The following example is illustrative.

Example 1 *Suppose that the seller arrives at date 1 for sure, while the buyer has an uncertain arrival time, distributed according to $G_B(t)$ with full support on all periods \mathbb{N} . Suppose that all payoff types are drawn i.i.d. in each period, with full support on $[\underline{\theta}_B, \bar{\theta}_B]$ for the buyer and full support on $[\underline{\theta}_S, \bar{\theta}_S]$ for the seller. If $\underline{\theta}_B \leq \underline{\theta}_S$ and $\bar{\theta}_B \leq \bar{\theta}_S$, then \bar{U}_1 defined in (12) is strictly negative for any blind mechanism satisfying E, BIC and BPC. Conversely, if either $\underline{\theta}_B > \underline{\theta}_S$ or $\bar{\theta}_B > \bar{\theta}_S$, then there*

exists $\bar{\delta}$ such that, for all $\delta \in (\bar{\delta}, 1)$, $\bar{U}_1 \geq 0$ for some blind mechanism satisfying these desiderata.

The result states that budget-balanced efficient trade is infeasible in case the lowest value type for the buyer and the highest cost type for the seller anticipate no gains from trade. In this case, $\bar{\Psi}_t$, as defined by (13), is equal to zero for all $t \geq 2$, while it is equal to $-\mathbb{E} \left[\left(\tilde{\theta}_{B,1} - \tilde{\theta}_{S,1} \right) x^E \left(\tilde{\theta}_{B,1}, \tilde{\theta}_{S,1} \right) \mid \tilde{\theta}_{B,1} \neq \emptyset \right]$ for $t = 1$. Conversely, if either $\underline{\theta}_B > \underline{\theta}_S$ or $\bar{\theta}_B > \bar{\theta}_S$, then gains from trade are anticipated with positive probability by either the lowest-value buyer or the highest-cost seller, and so $\bar{\Psi}_t > 0$ for all $t \geq 2$. We then have that \bar{U}_1 becomes non-negative for any δ close enough to one.

The environment of Example 1 is comparable to that in Athey and Miller (2007), where both agents are present from the beginning and the buyer and seller draw values and costs i.i.d. in each period from a common interval $[\underline{\theta}, \bar{\theta}]$ (with both the value and cost distributions having full support on this interval). While Athey and Miller show that efficient trade can be sustained with an ex-ante budget surplus provided the discount factor δ is at least one half, introducing uncertain arrival on just one side of the market renders this impossible for all $\delta \in (0, 1)$.¹⁹ Example 1 shows, however, that such a conclusion depends on the supports of buyer and seller values and costs. This dependence parallels observations in the literature on static trade with more than two traders: for instance, while Gresik and Satterthwaite (1989) observed that efficiency is unattainable irrespective of the number of traders when the support is common, Makowski and Mezzetti (1993) derived possibility results for a setting with at least two potential buyers and a support satisfying $\bar{\theta}_B > \bar{\theta}_S$.

4.2 Restriction on the supports of payoff types

We now comment on our support assumption; namely that, for all $\theta_{i,t-1} \in \Theta_{i,t-1}$, $[\underline{\theta}_{i,t}(\theta_{i,t-1}), \bar{\theta}_{i,t}(\theta_{i,t-1})] \subset \Theta_{i,t}$. This implies that, in any efficient mechanism, an agent who fails to participate at his arrival date, but who participates at the next

¹⁹Example 1 should also be compared to Gershkov, Moldovanu and Strack (2015) who analyze the limits to efficient implementation in a dynamic model where buyers arrive over time and where the planner learns through arriving buyers about the future arrival rate. While unrestricted subsidies can ensure an efficient allocation, limits on payments (in particular, a restriction to "winner pays") can jeopardize efficiency. Relatedly, Example 1 shows how limiting transfers by ruling out budget deficits jeopardizes efficiency in a setting where buyers arrive over time.

opportunity and then reports payoff types truthfully, induces efficient allocations from then on. This effectively – by payoff equivalence – ties the hands of the designer, nullifying any scope for punishing the deviation of delayed participation.²⁰ When our support restriction is not satisfied, budget-balanced efficient trade may be feasible even when the expression in (10) is negative as in the following example.

Example 2 *Suppose that the buyer's and seller's arrival dates are uncertain, with G_B and G_S each having full support on the dates \mathbb{N} . Suppose that, for each $i \in \{B, S\}$, each arrival date τ_i , and each $t > \tau_i$, $\theta_{i,t}$ is drawn i.i.d. from a non-degenerate and absolutely continuous distribution on $[\underline{\theta}, \bar{\theta}]$, a finite interval in \mathbb{R}_+ . Finally, suppose that, for each τ_B , θ_{B,τ_B} is smaller than $\underline{\theta}$, while for each τ_S , θ_{S,τ_S} is larger than $\bar{\theta}$, implying that surplus enhancing opportunities to trade do not exist at either the buyer's or seller's arrival dates. Then there exists a blind mechanism satisfying E, BB, BIC and BPC, and a public mechanism satisfying E, BB, PIC and PPC for any value $\delta \in (0, 1)$.*

To understand the result in Example 2, we explain only why EBS is satisfied by a blind mechanism satisfying E, BIC and BPC (leaving a proof for this example to the Appendix). Suppose the environment is blind, and consider a sequence of VCG mechanisms, one per period, with each agent earning the realized surplus from trade at each date. Now suppose that, to access these VCG mechanisms, each agent pays a fee on participation equal to the expected gains from trade from that date onwards. Paying this fee just once entitles the agent to participate forever after. Each agent then has an expected payoff from participation of exactly zero, so is willing to participate given that he can earn at most zero by participating in future. Hence the broker's ex-ante surplus is equal to the expected present value of all gains from trade.

Example 2 is "extreme" in the sense that, dropping the budget-balanced requirement, efficient trade can be implemented while ensuring agents earn zero expected rents. This is effectively achieved by "punishing" an agent who delays participation by denying the possibility of efficient trade until one period after participation (in particular, an agent who participates one period later incurs a one period delay in

²⁰The same would be true if our support restriction were not satisfied, but mechanisms were required to implement efficient allocations also "off path", i.e. for agents who have deviated in the past, say by failing to participate.

the first date he can trade, potentially denying an efficient trade). Two assumptions jointly permit the absence of buyer and seller rents in this example: (i) payoff types are drawn independently across periods, and (ii) the absence of opportunities for efficient trade upon arrival. However, the idea that weakening our support restriction permits the designer flexibility to punish delayed participation by imposing an inefficient trading rule applies more generally.

It is also worth noting the close parallel between violations of our support assumption and the possibility that agents are known not to arrive at certain dates (as discussed in Section 4.1). In both cases, agent rents can be reduced by committing to an inefficient policy in case of a deviation. For instance, when one of the agents is known to arrive at date one, a commitment can be made not to permit any trade if the agent fails to participate at that date (we discuss relaxing such commitments in Section 4.4 below).

4.3 Constrained-efficient mechanisms

A natural question, which turns out to be intimately related to the preceding discussion on payoff-type supports, is constrained-efficiency. For instance, in environments where any efficient mechanism eliciting truthful reporting runs an expected budget deficit, what allocation $(x_t(\cdot, \cdot))_{t \in \mathbb{N}}$ maximizes the ex-ante surplus

$$\mathbb{E} \left[\sum_{t=\max\{\tilde{\tau}_S, \tilde{\tau}_B\}}^{\infty} \delta^{t-1} \left(\tilde{\theta}_{B,t} - \tilde{\theta}_{S,t} \right) x_t \left(\tilde{\theta}_{B,\tau_B}^t, \tilde{\theta}_{S,\tau_S}^t \right) \right] \quad (15)$$

subject to either EBS or BB?

A preliminary observation is that, at least when payoff types are not too persistent (in a sense formalized by the "impulse responses" introduced by Pavan, Segal and Toikka, 2014), efficient allocations can be implemented at dates long after contracting without ceding large information rents to agents who have favorable initial information (say, to a buyer with a high initial value, or a seller with a low initial cost).²¹ We should therefore anticipate that the fraction of surplus foregone by a

²¹For instance, the mechanism may specify no trade for any agent until at least m periods since participation, and efficient trade thereafter (provided the other agent has also participated for at least m periods). Such a mechanism can be found that is incentive compatible and lowers agent rents relative to a mechanism stipulating efficient trade at all dates.

constrained-efficient mechanism (as envisioned above) would often vanish as the discount factor δ approaches one. This is not to say that dynamic arrivals do not lead to large reductions in the fraction of surplus attainable for fixed values of δ .

While we do not provide a full exploration of constrained-efficient mechanisms, a simple example will be suggestive of the forces at play. To this end, suppose that, for each agent i , in each period t from the arrival date τ_i onwards, agent i makes an i.i.d. draw of $\theta_{i,t}$ from an absolutely continuous distribution F_i with support on $[\underline{\theta}_i, \bar{\theta}_i]$. Suppose these supports overlap: i.e., $\bar{\theta}_B > \underline{\theta}_S$ and $\underline{\theta}_B < \bar{\theta}_S$. Following the same steps as in Garrett (2016), one can deduce that the smallest expected rents in an incentive-compatible mechanism with allocation rule $(x_t(\cdot, \cdot))_{t \in \mathbb{N}}$ equal

$$\mathbb{E} \left[\delta^{\tilde{\tau}_B - 1} \frac{G_B(\tilde{\tau}_B)}{g_B(\tilde{\tau}_B)} \frac{1 - F_B(\tilde{\theta}_{B, \tilde{\tau}_B})}{f_B(\tilde{\theta}_{B, \tilde{\tau}_B})} x_{\tilde{\tau}_B}(\tilde{\theta}_{B, \tilde{\tau}_B}, \tilde{\theta}_{S, \tilde{\tau}_B}^{\tilde{\tau}_B}) \right] \quad (16)$$

for the buyer and

$$\mathbb{E} \left[\delta^{\tilde{\tau}_S - 1} \frac{G_S(\tilde{\tau}_S)}{g_S(\tilde{\tau}_S)} \frac{F_S(\tilde{\theta}_{S, \tilde{\tau}_S})}{f_S(\tilde{\theta}_{S, \tilde{\tau}_S})} x_{\tilde{\tau}_S}(\tilde{\theta}_{B, \tilde{\tau}_B}^{\tilde{\tau}_S}, \tilde{\theta}_{S, \tilde{\tau}_S}) \right] \quad (17)$$

for the seller.

To understand these expressions, note that (in a mechanism that minimizes agent rents) an agent i 's expected information rents if arriving at τ_i do not depend on allocations for dates $t > \tau_i$. This follows because payoff types are independent across time, and so agents lack private information about future payoff types at the time of contracting. Consider the buyer's case. If the arrival date τ_B were certain and commonly known (with the seller's arrival date still uncertain), then the only source of information rents would be the date τ_B value θ_{B, τ_B} , and the buyer's ex-ante expected rents would hence equal

$$\mathbb{E} \left[\delta^{\tau_B - 1} \frac{1 - F_B(\tilde{\theta}_{B, \tau_B})}{f_B(\tilde{\theta}_{B, \tau_B})} x_{\tau_B}(\tilde{\theta}_{B, \tau_B}, \tilde{\theta}_{S, \tau_B}^{\tau_B}) \right],$$

which is the standard expression from static mechanism design. The expression (16), however, is more than simply its expectation over uncertain realizations of τ_B . In

particular, when the buyer's arrival time is uncertain, inducing participation at all dates before some participation date τ_B requires the same rent to be attributed to the buyer if arriving at any earlier date, an event that has probability $G_B(\tau_B - 1)$. This explains why the ratio $\frac{G_B(\tilde{\tau}_B)}{g_B(\tilde{\tau}_B)}$ (which captures the likelihood of earlier arrival relative to arrival at date $\tilde{\tau}_B$) appears in expression (16). An analogous observation holds for the seller.

The requirement that a broker financing trade breaks even in expectation (i.e., constraint EBS) can then be written as

$$\begin{aligned} & \mathbb{E} \left[\begin{aligned} & \delta^{\tilde{\tau}_B-1} \frac{G_B(\tilde{\tau}_B)}{g_B(\tilde{\tau}_B)} \frac{1-F_B(\tilde{\theta}_{B,\tilde{\tau}_B})}{f_B(\tilde{\theta}_{B,\tilde{\tau}_B})} x_{\tilde{\tau}_B} \left(\tilde{\theta}_{B,\tilde{\tau}_B}, \tilde{\theta}_{S,\tilde{\tau}_B} \right) \\ & + \delta^{\tilde{\tau}_S-1} \frac{G_S(\tilde{\tau}_S)}{g_S(\tilde{\tau}_S)} \frac{F_S(\tilde{\theta}_{S,\tilde{\tau}_S})}{f_S(\tilde{\theta}_{S,\tilde{\tau}_S})} x_{\tilde{\tau}_S} \left(\tilde{\theta}_{B,\tilde{\tau}_B}, \tilde{\theta}_{S,\tilde{\tau}_S} \right) \end{aligned} \right] \\ \leq & \mathbb{E} \left[\sum_{t=\max\{\tilde{\tau}_B,\tilde{\tau}_S\}}^{\infty} \delta^{t-1} \left(\tilde{\theta}_{B,t} - \tilde{\theta}_{S,t} \right) x_t \left(\tilde{\theta}_{B,\tilde{\tau}_B}^t, \tilde{\theta}_{S,\tilde{\tau}_S}^t \right) \right], \end{aligned} \quad (18)$$

which states that the expected surplus from trade covers expected rents. A first pass at maximizing efficiency subject to the broker breaking even would then be to maximize (15) subject to (18). The solution to this "relaxed program" coincides with the solution to the problem of interest provided that $\frac{1-F_B(\cdot)}{f_B(\cdot)}$ is non-increasing and $\frac{F_S(\cdot)}{f_S(\cdot)}$ is non-decreasing (which will ensure monotonicity of the allocation in each agent's payoff type).

Example 3 Suppose that each agent i 's payoff types are drawn i.i.d. in each period from F_i (as defined above), that the supports overlap, and that $\frac{1-F_B(\cdot)}{f_B(\cdot)}$ is non-increasing while $\frac{F_S(\cdot)}{f_S(\cdot)}$ is non-decreasing. Consider the allocation rule $(x_i^*(\cdot, \cdot))_{t \in \mathbb{N}}$ that maximizes ex-ante surplus (15) in a blind mechanism satisfying BIC, BPC and EBS. This sets $x_i^*(\theta_{B,\tau_B}^t, \theta_{S,\tau_S}^t)$ equal to one if

$$\theta_{B,t} - \theta_{S,t} \geq \begin{cases} \frac{\lambda}{1+\lambda} \left[\frac{G_B(\tau_B)}{g_B(\tau_B)} \frac{1-F_B(\theta_{B,\tau_B})}{f_B(\theta_{B,\tau_B})} + \frac{G_S(\tau_S)}{g_S(\tau_S)} \frac{F_S(\theta_{S,\tau_S})}{f_S(\theta_{S,\tau_S})} \right] & \text{for } t = \tau_B = \tau_S \\ \frac{\lambda}{1+\lambda} \frac{G_B(\tau_B)}{g_B(\tau_B)} \frac{1-F_B(\theta_{B,\tau_B})}{f_B(\theta_{B,\tau_B})} & \text{for } t = \tau_B > \tau_S \\ \frac{\lambda}{1+\lambda} \frac{G_S(\tau_S)}{g_S(\tau_S)} \frac{F_S(\theta_{S,\tau_S})}{f_S(\theta_{S,\tau_S})} & \text{for } t = \tau_S > \tau_B \\ 0 & \text{for } t > \max\{\tau_B, \tau_S\} \end{cases}, \quad (19)$$

where $\lambda > 0$ is the Lagrange multiplier on (18), and zero otherwise. The same

allocation is part of a public mechanism that satisfies PIC, PPC and EBS, which is hence constrained-efficient among public mechanisms satisfying these desiderata.

A few comments are in order. First, because payoff types are not persistent over time, allocations are efficient as soon as both agents have participated in the past (the example is a limiting case of the settings with limited type persistence, as discussed above). Allocations are more distorted (i.e., more trade that is efficient fails to occur) on a date t if both agents arrive at that date (i.e., if $t = \tau_B = \tau_S$) than if only one agent arrives at that date while the other has already arrived (i.e., if $t = \tau_B > \tau_S$ or $t = \tau_S > \tau_B$). Since distorted trade lasts exactly one period, with efficient trade occurring thereafter, the lost surplus due to the constraint EBS shrinks to zero as a fraction of the total efficient surplus as δ approaches one.

While trade is distorted only on the first date it can occur (i.e., the first date at which both agents have arrived), the size of distortions depend both on the timing of this event and the distributions G_B and G_S of agent arrival times. A natural possibility is that $\frac{G_i(t)}{g_i(t)}$ is increasing in t , which is the case for instance if each G_i is the geometric distribution with parameter $\mu_i \in (0, 1)$ (in which case $G_i(t) = 1 - (1 - \mu_i)^t$ for each i). In such cases, distortions are larger if the first date where trade is possible (i.e., $\max\{\tau_B, \tau_S\}$) occurs later. This reflects that distortions introduced at a given date reduce the rents that must be left to agents arriving at all earlier dates, and that the probability of such earlier arrival necessarily increases with time. Under our maintained assumption that G_i has full support on the set of all dates \mathbb{N} , notice that $\frac{G_i(t)}{g_i(t)}$ necessarily grows without bound. After enough time, the first trading date is hence characterized by large distortions: trade only takes place at this date if either the buyer's value is very close to $\bar{\theta}_B$ or the seller's cost is very close to $\underline{\theta}_S$ (both are required if $\tau_B = \tau_S$).

Finally, note that the result in Example 3 considers the constraint that the broker must expect a surplus (i.e., condition EBS). In the blind environment, it is possible to take the constrained-efficient mechanism satisfying BIC, BPC and EBS, and adjust transfers so that it also satisfies BB. For instance, following the same approach as for AGV mechanisms (as in d'Aspremont and Gérard-Varet, 1979), it is possible to adjust transfers period by period so that BB holds while BIC continues to be satisfied. As in the proof of Proposition 2, one can then make "dynamic" adjustments to transfers that redistribute surplus across time and between players to ensure BPC. However, when the mechanism is public, analogous steps are often not possible; i.e.,

the allocation $(x_t^*(\cdot, \cdot))_{t \in \mathbb{N}}$ of Example 3 cannot be implemented by a mechanism that satisfies BB, PIC and PPC. For instance, when $G_i(t)$ is the geometric distribution for each i , we see that

$$\begin{aligned} & \mathbb{E} \left[\begin{aligned} & \delta^{\tilde{\tau}_B - 1} \frac{G_B(\tilde{\tau}_B)}{g_B(\tilde{\tau}_B)} \frac{1 - F_B(\tilde{\theta}_{B, \tilde{\tau}_B})}{f_B(\tilde{\theta}_{B, \tilde{\tau}_B})} x_{\tilde{\tau}_B}^* \left(\tilde{\theta}_{B, \tilde{\tau}_B}, \tilde{\theta}_{S, \tilde{\tau}_B} \right) \\ & + \delta^{\tilde{\tau}_S - 1} \frac{G_S(\tilde{\tau}_S)}{g_S(\tilde{\tau}_S)} \frac{F_S(\tilde{\theta}_{S, \tilde{\tau}_S})}{f_S(\tilde{\theta}_{S, \tilde{\tau}_S})} x_{\tilde{\tau}_S}^* \left(\tilde{\theta}_{B, \tilde{\tau}_B}, \tilde{\theta}_{S, \tilde{\tau}_S} \right) \end{aligned} \mid \tilde{\theta}_{B,1} = \tilde{\theta}_{S,1} = \emptyset \right] \\ < \mathbb{E} \left[\sum_{t=\max\{\tilde{\tau}_B, \tilde{\tau}_S\}}^{\infty} \delta^{t-1} \left(\tilde{\theta}_{B,t} - \tilde{\theta}_{S,t} \right) x_t^* \left(\tilde{\theta}_{B,t}, \tilde{\theta}_{S,t} \right) \mid \tilde{\theta}_{B,1} = \tilde{\theta}_{S,1} = \emptyset \right]. \quad (20) \end{aligned}$$

The follows by a comparison to constraint (18), which holds with equality for the allocation $(x_t^*(\cdot, \cdot))_{t \in \mathbb{N}}$. In particular, conditional on no arrival at date 1, the future evolution of arrivals, values and costs from date 2 onwards is the same as at the beginning of the game, yet future trade is more distorted. The inequality (20) implies that the allocation $(x_t^*(\cdot, \cdot))_{t \geq 2}$ can be implemented by a PIC and PPC mechanism that generates a strict budget surplus from date 2 onwards, conditional on the event that no agent arrives at date 1. Alternatively, since budget balance requires the surplus to be fully distributed between the agents, any budget balanced mechanism must yield continuation payoffs larger than the smallest possible value in any PIC and PPC mechanism implementing $(x_t^*(\cdot, \cdot))_{t \in \mathbb{N}}$. This implies that a broker in this arrangement must expect an ex-ante loss, and hence that budget-balanced trade is infeasible.

4.4 Limited designer commitment

Our results shed light on a strong form of commitment relied on in the existing possibility results for efficient repeated bilateral trade (reviewed above). In particular, we have noted that the assumption that agent arrival times are commonly known permits a designer with commitment ability to exclude agents who fail to participate on these dates. We have shown that this commitment can be crucial for reducing agent rents, and hence obtaining criteria such as EBS or BB subject to incentive and participation constraints. However, commitments to exclude may not be credible in some allocation problems. Any central authority charged with designing a mechanism for efficient outcomes would presumably face pressure to still implement efficient outcomes in case one of the trading partners failed to participate in the cho-

sen mechanism at the anticipated date. In a decentralized environment, we might expect potential trading partners to renegotiate, at least to prevent the complete break down of trade, if one of these partners turned up to the bargaining table later than expected.

To see the implications of a lack of commitment more concretely, suppose that agents draw values and costs i.i.d. in each period from a commonly known arrival date (say $t = 1$), and suppose that these distributions have common support. While the ability to exclude non-participants permits efficient and budget-balanced trade for any $\delta \geq 1/2$, such trade is infeasible for any δ if the mechanism instead implements efficiency from the first date at which both agents have participated (no matter when that happens to be).²²

This observation contrasts with what is seen in many dynamic mechanism design problems with the objective of efficiency. For instance, Bergemann and Valimaki (2010, p 772) note that: "the dynamic pivot mechanism is time-consistent and the social choice function can be implemented by a sequential mechanism without any ex ante commitment by the designer". Yet, if there are constraints on transfers (such as EBS or BB), we have seen that this is not a general property of efficient mechanisms.

4.5 Other allocation problems

Our main characterization results can be easily extended to many other allocation problems where agents have quasi-linear payoffs, such as public goods problems and trade involving more than two agents. For instance, as in Proposition 1, if agents' arrival times have full support across all dates and are their private information, the maximum expected broker surplus among efficient mechanisms satisfying incentive compatibility and participation constraints is attainable by a sequence of appropriately chosen VCG mechanisms. In turn, this permits a characterization of environments in which efficient allocations are feasible satisfying revenue neutrality or budget balance constraints. Again, the relevant condition is a "weighted average" of the condition for static problems. Here, it is worth pointing out that, in various allocation problems, such as settings with more than two traders, efficiency can be attainable in

²²Here, we use that the designer continues to believe that values and costs are drawn from the same common support in each period, even after a deviation to non-participation. While this makes sense in light of i.i.d. draws, the question of designer beliefs is more subtle when values and costs are correlated over time.

the static problem on a given date, subject to budget balance and relevant incentive constraints, even if the supports of payoff types overlap. A condition analogous to (10) can be used to distinguish the feasibility of efficient trade in such problems.

4.6 Allocation-dependent processes

Another possible extension is to permit the evolution of payoff types to depend on past allocations or actions; see Athey and Segal (2013) for a model where this is permitted. While care would be needed to adapt our results to this case, it seems reasonable to conjecture that budget-balanced trade can become possible even in instances when the distributions of values and costs overlap. In particular, if trade of the good itself tends to increase the buyer's future values and lower the seller's future costs, then these stochastic "improvements" could act as a reward for early participation in the mechanism, relaxing participation constraints.

5 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 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), or if the designer cannot commit to shutting down trade in the event of late participation, then efficient budget-neutral trade can be (much) more difficult to sustain. In the classic bilateral trade problem, overlapping supports for buyer and seller values is enough to render budget-balanced and efficient trade infeasible.

References

- [1] Athey, Susan and David A. Miller (2007), ‘Efficiency in repeated trade with hidden valuations,’ *Theoretical Economics*, 2, 299-354.
- [2] Athey, Susan and Ilya Segal (2007), ‘Designing Dynamic Mechanisms for Dynamic Bilateral Trading Games,’ *American Economic Review*, 97, 131-136.
- [3] Athey, Susan and Ilya Segal (2013), ‘An Efficient Dynamic Mechanism,’ *Econometrica*, 81, 2463-2485.
- [4] Baron, David P., and David Besanko (1984), ‘Regulation and Information in a Continuing Relationship,’ *Information Economics and Policy*, 1(3), 267–302.
- [5] Battaglini, Marco (2005), ‘Long-Term Contracting with Markovian Consumers,’ *American Economic Review*, 95, 637-658.
- [6] Battaglini, Marco and Rohit Lamba (2015), ‘Optimal Dynamic Contracting: the First-Order Approach and Beyond,’ mimeo Princeton University and Pennsylvania State University.
- [7] Bergemann, Dirk and Juuso Valimaki (2010), ‘The Dynamic Pivot Mechanism,’ *Econometrica*, 78, 771-789.
- [8] Besanko, David (1985), ‘Multi-Period Contracts between Principal and Agent with Adverse Selection,’ *Economics Letters*, 17, 33-37.
- [9] Board, Simon and Andrzej Skrzypacz (2016), ‘Revenue Management with Forward-Looking Buyers,’ *Journal of Political Economy*, 124, 1046-1087.
- [10] Boleslavsky, Raphael and Maher Said (2013), ‘Progressive Screening: Long-Term Contracting with a Privately Known Stochastic Process,’ *Review of Economic Studies*, 80, 1-34.
- [11] Chatterjee, Kalyan and William Samuelson (1983), ‘Bargaining under Incomplete Information,’ *Operations Research*, 31, 835-851.
- [12] Courty, Pascal. and Li Hao (2000), ‘Sequential Screening,’ *Review of Economic Studies*, 67, 697–717.

- [13] Deb, Rahul and Maher Said (2015), ‘Dynamic screening with limited commitment,’ *Journal of Economic Theory*, 159B, 891-928.
- [14] Ely, Jeffrey C., Daniel Garrett and Toomas Hinnosaar (forthcoming), ‘Overbooking,’ *Journal of the European Economic Association*.
- [15] Eso, Peter and Balazs Szentes (2007), ‘Optimal Information Disclosure in Auctions and the Handicap Auction,’ *Review of Economic Studies*, 74(3), 705-731.
- [16] Eso, Peter and Balazs Szentes (forthcoming), ‘Dynamic Contracting: An Irrelevance Result,’ *Theoretical Economics*.
- [17] Garratt, Rod and Marek Pycia (2016), ‘Efficient Bilateral Trade,’ mimeo UC Santa Barbara and UCLA.
- [18] Garrett, Daniel (2016), ‘Dynamic Mechanism Design: Dynamic Arrivals and Changing Values,’ mimeo Toulouse School of Economics.
- [19] Garrett, Daniel (forthcoming), ‘Intertemporal Price Discrimination: Dynamic Arrivals and Changing Values,’ *American Economic Review*.
- [20] Gershkov, Alex, Benny Moldovanu and Philipp Strack (2015), ‘Efficient Dynamic Allocation with Strategic Arrivals,’ mimeo Hebrew University of Jerusalem, University of Surrey, University of Bonn and UC Berkeley.
- [21] Gershkov, Alex, Benny Moldovanu and Philipp Strack (2016), ‘Revenue Maximizing Mechanisms with Strategic Customers and Unknown Demand,’ mimeo Hebrew University of Jerusalem, University of Surrey, University of Bonn and UC Berkeley.
- [22] Gresik, Thomas A. and Mark A. Satterthwaite (1989), ‘The rate at which a simple market converges to efficiency as the number of traders increases: An asymptotic result for optimal trading mechanisms,’ *Journal of Economic Theory*, 48, 304-332.
- [23] Kräbmer, Daniel and Roland Strausz (2015), ‘Optimal Sales Contracts with Withdrawal Rights,’ *Review of Economic Studies*, 82, 762-790.
- [24] Krishna, Vijay and Motty Perry (1998), ‘Efficient Mechanism Design,’ mimeo Hebrew University of Jerusalem and Pennsylvania State University.

- [25] Lamba, Rohit (2013), ‘Repeated Bargaining: A Mechanism Design Approach,’ mimeo Princeton University and Pennsylvania State University.
- [26] Makowski, Louis and Claudio Mezzetti (1993), ‘The Possibility of Efficient Mechanisms for Trading an Indivisible Object,’ *Journal of Economic Theory*, 59, 451-465.
- [27] Makowski, Louis and Claudio Mezzetti (1994), ‘Bayesian and Weakly Robust First Best Mechanisms: Characterizations,’ *Journal of Economic Theory*, 64, 500-519.
- [28] Myerson, Roger B. (1986), ‘Multistage Games with Communication,’ *Econometrica*, 54, 323-358.
- [29] Myerson, Roger B. and Mark A. Satterthwaite (1983), ‘Efficient mechanisms for bilateral trade,’ *Journal of Economic Theory*, 29, 265-281.
- [30] Pavan, Alessandro, Ilya Segal and Juuso Toikka (2014), ‘Dynamic Mechanism Design: a Myersonian Approach,’ *Econometrica*, 82, 601-653.
- [31] Simon, Herbert A. (1955), ‘A Behavioral Model of Rational Choice,’ *Quarterly Journal of Economics*, 69, 99-118.
- [32] Skrzypacz, Andrzej and Juuso Toikka (2013), ‘Mechanisms for Repeated Trade,’ *American Economic Journal: Microeconomics*, 7, 252-293.
- [33] Williams, Steven R. (1999), ‘A characterization of efficient, Bayesian incentive compatible mechanisms,’ *Economic Theory*, 14, 155-180.
- [34] Wolitzky, Alexander (forthcoming), ‘Mechanism Design with Maxmin Agents: Theory and an Application to Bilateral Trade,’ *Theoretical Economics*.
- [35] Yoon, Kiho (2015), ‘On budget balance of the dynamic pivot mechanism,’ *Games and Economic Behavior*, 94, 206-213.

Appendix: Proofs of results

Proof of Lemma 1. Given our assumption on the evolution of payoff types and focus on BIC mechanisms, a necessary condition for agent i to be willing to participate at date τ_i when his date- τ_i type is θ_{i,τ_i} is given by (4). This implies $\inf_{\theta_{i,\tau_i} \in \Theta_{i,\tau_i}} \left\{ V_{i,\tau_i}^\Omega(\theta_{i,\tau_i}) - \delta \mathbb{E} \left[V_{i,\tau_i+1}^\Omega(\tilde{\theta}_{i,\tau_i+1}) \mid \theta_{i,\tau_i} \right] \right\} \geq 0$ in any blind mechanism satisfying BIC and BPC. If the expression is strictly positive, for any i and τ_i , we can increase the broker's surplus U_1 simply by reducing $p_{i,\tau_i}(\theta_{i,\tau_i}, \theta_{-i,\tau_i}^{\tau_i})$ by a uniform constant smaller than $\inf_{\theta_{i,\tau_i} \in \Theta_{i,\tau_i}} \left\{ V_{i,\tau_i}^\Omega(\theta_{i,\tau_i}) - \delta \mathbb{E} \left[V_{i,\tau_i+1}^\Omega(\tilde{\theta}_{i,\tau_i+1}) \mid \theta_{i,\tau_i} \right] \right\}$. This reduces $V_{i,\tau_i}^\Omega(\theta_{i,\tau_i})$ by the same amount, relaxing the constraint (4) for agent i at $\tau_i - 1$, but without violating it at any other date. In particular, if the original mechanism satisfied BIC and BPC, then the adjusted mechanism satisfies these conditions as well. ■

Proof of Proposition 1. Consider any blind mechanism satisfying E, BIC, BPC and the condition of Lemma 1 for all dates. Given our continuity assumption on $F_t^{i,Tr}$, there exists for each i and t a $\theta_{i,t}^* \in \Theta_{i,t}$ satisfying

$$V_{i,t}^\Omega(\theta_{i,t}^*) = \delta \left[V_{i,t+1}^\Omega(\tilde{\theta}_{i,t+1}) \mid \tilde{\theta}_{i,t} = \theta_{i,t}^* \right].$$

Hence, for each i and t ,

$$V_{i,t}^\Omega(\theta_{i,t}^*) = \delta V_{i,t+1}^\Omega(\theta_{i,t+1}^*) + \delta \mathbb{E} \left[\left(V_{i,t+1}^\Omega(\tilde{\theta}_{i,t+1}) - V_{i,t+1}^\Omega(\theta_{i,t+1}^*) \right) \mid \tilde{\theta}_{i,t} = \theta_{i,t}^* \right].$$

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

$$V_{i,t}^\Omega(\theta_{i,t}^*) = \sum_{s=1}^l \delta^s \mathbb{E} \left[\left(V_{i,t+s}^\Omega(\tilde{\theta}_{i,t+s}) - V_{i,t+s}^\Omega(\theta_{i,t+s}^*) \right) \mid \tilde{\theta}_{i,t+s-1} = \theta_{i,t+s-1}^* \right] + \delta^l V_{i,t+l}^\Omega(\theta_{i,t+l}^*).$$

Since the option is always available to never participate, we must have $V_{i,t+l}^\Omega(\theta_{i,t+l}^*) \geq 0$ for all l . This, together with the assumption of bounded payoff types, implies that $\lim_{l \rightarrow +\infty} \delta^l V_{i,t+l}^\Omega(\theta_{i,t+l}^*) = 0$ in any mechanism that maximizes U_1 . Hence, the smallest feasible value of $V_{i,t}^\Omega(\theta_{i,t}^*)$ is

$$\sum_{s=1}^{\infty} \delta^s \mathbb{E} \left[\left(V_{i,t+s}^\Omega(\tilde{\theta}_{i,t+s}) - V_{i,t+s}^\Omega(\theta_{i,t+s}^*) \right) \mid \tilde{\theta}_{i,t+s-1} = \theta_{i,t+s-1}^* \right].$$

By the payoff equivalence property of Assumption 1, $V_{i,t+s}^\Omega(\theta_{i,t+s}) - V_{i,t+s}^\Omega(\theta_{i,t+s}^*)$ is the same across all blind mechanisms satisfying BIC and E, and this is true for all $s \geq 1$. Thus, any mechanism that is BIC and E, and which satisfies the condition of Lemma 1 at all dates as well as $\lim_{t \rightarrow +\infty} \delta^t V_{i,t+l}^\Omega(\theta_{i,t+l}^*) = 0$, has the same value of U_1 . One example of such a mechanism is the sequence of VCG-* mechanisms specified in the proposition. Such mechanisms are BIC because the static VCG mechanism is incentive compatible, and the mechanism an agent faces at future dates does not depend on current reports. Each agent is also guaranteed a non-negative payoff from participating in a given date- t mechanism, while retaining the right to participate in the same future mechanisms he would engage in if delaying participation. So the mechanism satisfies BPC. 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., $\theta_{B,t}^* = \underline{\theta}_{B,t}$ and $\theta_{S,t}^* = \bar{\theta}_{S,t}$). Hence, the condition of Lemma 1 is satisfied. Also, using that payoff types are bounded, $\lim_{t \rightarrow +\infty} \delta^{t-1} V_{i,t}^\Omega(\theta_{i,t}^*) = 0$ for each i . Hence agent rents are as small as possible. Hence U_1 is as large as possible in a blind mechanism satisfying E, BIC, BPC and NPS.

For the public environment, consider a sequence of efficient static mechanisms with payments given, if both agents have arrived by date t , by

$$\begin{aligned} p_{B,t}^{VCG^{*-}P}(\theta_{B,\tau_B}^t, \theta_{S,\tau_S}^t) &= -\theta_{S,t} x^E(\theta_{B,t}, \theta_{S,t}) \\ &\quad - \mathbb{E} \left[\left(\underline{\theta}_{B,t} - \tilde{\theta}_{S,t} \right) x^E(\underline{\theta}_{B,t}, \tilde{\theta}_{S,t}) \mid \tilde{\theta}_{S,\tau_S}^{t-1} = \theta_{S,\tau_S}^{t-1}, \tilde{\theta}_{S,t} \neq \emptyset \right] \end{aligned}$$

for the buyer, and

$$\begin{aligned} p_{S,t}^{VCG^{*-}P}(\theta_{S,\tau_S}^t, \theta_{B,\tau_B}^t) &= \theta_{B,t} x^E(\theta_{B,t}, \theta_{S,t}) \\ &\quad - \mathbb{E} \left[\left(\tilde{\theta}_{B,t} - \bar{\theta}_{S,t} \right) x^E(\tilde{\theta}_{B,t}, \bar{\theta}_{S,t}) \mid \tilde{\theta}_{B,\tau_B}^{t-1} = \theta_{B,\tau_B}^{t-1}, \tilde{\theta}_{B,t} \neq \emptyset \right] \end{aligned}$$

for the seller. If at least one of the agents has not arrived by date t , then no payments are made. Again each static mechanism is incentive compatible, and an agent's reports do not affect his payments in future (static) mechanisms, so the sequence of static mechanisms is PIC. Agents earn non-negative payoffs by participating in each static mechanism (and such participation does not impose future obligations) so the sequence of static mechanisms satisfies PPC. Finally, note that, by the law of iterated

expectations, the ex-ante expectation of payments are identical under the mechanism for the blind environment as for the one specified here for the public environment. Hence the broker's expected surplus U_1 is identical for both mechanisms (and hence is the highest surplus achievable in the public environment). ■

Proof of Proposition 2. The expression (10) for U_1 follows from the arguments in the main text. As we have noted, budget-balanced efficient trade is infeasible if $U_1 < 0$.

If $U_1 \geq 0$, then we can construct a blind mechanism satisfying BB, E, BIC and BPC as follows. Recall that a date- t budget-balanced, efficient and incentive-compatible static mechanism $\Omega_t^{static} = \langle x^E, p_{B,t}, p_{S,t} \rangle$ is such that

$$W_{B,t}^{\Omega_t^{static}}(\underline{\theta}_{B,t}) + W_{S,t}^{\Omega_t^{static}}(\bar{\theta}_{S,t}) = \Psi_t,$$

with Ψ_t given by (11) and where $W_{i,t}^{\Omega_t^{static}}(\theta_{i,t})$ gives the expected payoff from truthful reporting in the static mechanism conditional on both agents having arrived by t , and on i 's date- t payoff type $\theta_{i,t}$.²³ This observation is standard and can be understood from payoff equivalence to the static VCG mechanism. An example of such a static mechanism is the AGV mechanism, following d'Aspremont and Gérard-Varet (1979).

Given such a budget-balanced mechanism, it is then possible to construct a sequence of budget-balanced static mechanisms $\Omega_t^\#$ such that expected payoffs from participating (with value $\underline{\theta}_{B,t}$ for the buyer and $\bar{\theta}_{S,t}$ for the seller, and assuming the other agent participates) are given by

$$\begin{aligned} W_{B,t}^{\Omega_t^\#}(\underline{\theta}_{B,t}) &= \min\{0, \Psi_t\} \\ W_{S,t}^{\Omega_t^\#}(\bar{\theta}_{S,t}) &= \max\{0, \Psi_t\}. \end{aligned}$$

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 $\Omega_t^\#$ is $G_S(t) \min\{0, \Psi_t\}$ when his value for the good is $\underline{\theta}_{B,t}$, while the seller's is $G_B(t) \max\{0, \Psi_t\}$ when his cost is $\bar{\theta}_{S,t}$.

²³Note that we are only conditioning on agents having arrived by date t , and hence view all agents who have arrived by date t as participating in Ω_t^{static} .

We can then modify the sequence of static mechanisms $(\Omega_t^\#)_{t \geq 1}$ 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. If $\Psi_1 < 0$, then in the first period t such that $\Psi_t > 0$, require the seller (if participating) to pay the buyer either $\frac{G_S(1)|\Psi_1|}{\delta^{t-1}G_S(t)}$, or $\frac{G_B(t)\Psi_t}{G_B(1)}$ if the latter is smaller. In the first case, the buyer with payoff type $\underline{\theta}_{B,1}$ is indifferent to participating at date 1, while in the latter the seller with payoff type $\bar{\theta}_{S,t}$ is indifferent to participating (since he makes a payment $\frac{G_B(t)\Psi_t}{G_B(1)}$ with probability $G_B(1)$). In the latter case, we continue to the next date $t' > t$ such that $\Psi_{t'} > 0$ and require a payment to the date-1 participating buyer of either $\frac{G_S(1)|\Psi_1| - \frac{\Psi_t G_B(t) G_S(t) \delta^{t-1}}{G_B(1)}}{\delta^{t'-1} G_S(t')}$, or $\frac{G_B(t')\Psi_{t'}}{G_B(1)}$ if the latter is smaller. Continuing this way, we can ensure that the expected loss of the buyer arriving at date 1 with type $\underline{\theta}_{B,1}$, i.e. $G_S(1)|\Psi_1|$, is paid for by the seller (in expectation), if

$$G_S(1)|\Psi_1| \leq \sum_{t=2}^{\infty} \delta^{t-1} \max \left\{ \frac{G_B(t)\Psi_t}{G_B(1)}, 0 \right\} G_S(t),$$

or equivalently

$$G_B(1)G_S(1)|\Psi_1| \leq \sum_{t=2}^{\infty} \delta^{t-1} \max \{G_B(t)G_S(t)\Psi_t, 0\},$$

which is true because $U_1 \geq 0$.

We can then proceed sequentially by ensuring participation of buyers in every date- t mechanism. At a generic date t , we specify payments from the seller starting with the first date s at which the highest-cost seller, $\bar{\theta}_{S,s}$, expects positive rent from participating in the date- s mechanism. Ensuring participation of buyers up to date t in a budget-balanced mechanism (while maintaining seller participation at all dates) is then feasible provided that, for each $s \leq t$,

$$-\delta^{s-1}G_B(s)G_S(s) \min \{\Psi_s, 0\} \leq y_s$$

for $(y_s)_{s=1}^t$ satisfying

$$\sum_{s=1}^t y_s \leq \sum_{s=1}^{\infty} \delta^{s-1} \max \{G_B(s) G_S(s) \Psi_s, 0\}.$$

This is guaranteed for all t by $U_1 \geq 0$. Hence, participation can be guaranteed at all dates while maintaining budget balance.

It remains to check that the adjusted mechanism satisfies NPS. Note first that all payments given by $(\Omega_t^\#)_{t \geq 1}$ are uniformly bounded across realizations of agents' information, because payoff types are uniformly bounded. Hence their discounted sum is uniformly bounded. Second, note that the additional payments *from* the seller are always positive and have expected present value no greater than

$$\sum_{t=1}^{\infty} \delta^{t-1} \max \{G_B(t) G_S(t) \Psi_t, 0\} < +\infty.$$

Also, the additional payments *to* the buyer are always positive, with the expected present value equal to

$$\sum_{t=1}^{\infty} \delta^{t-1} |\min \{G_B(t) G_S(t) \Psi_t, 0\}| < +\infty.$$

These observations together imply that NPS is satisfied.

The final part of the result, i.e. the sufficient condition for $U_1 < 0$, follows from arguments in the main text. ■

Proof of Proposition 3. That the broker's expected surplus \bar{U}_1 is given by (12) in any broker-optimal mechanism implementing efficient trade follows using payoff equivalence to the sequence of VCG-* mechanisms described in the main text (the proof follows closely those of Propositions 1 and 2, and is hence omitted). It is then immediate that budget-balanced efficient trade is infeasible when $\bar{U}_1 < 0$.

In case $\bar{U}_1 \geq 0$ we construct a public mechanism satisfying BB, E, PIC and PPC, from which follows the existence of a blind mechanism satisfying BB, E, BIC and BPC.

Suppose first that (i) the seller reports his costs in each period from date 1, and (ii) if the buyer participates at τ_B , then a "balanced team mechanism", as described by Athey and Segal (2013), is played from τ_B onwards (calculated by taking the

seller's cost at $\tau_B - 1$ to be the reported one). This defines a dynamic mechanism in which the seller participates from date 1 and the buyer participates from his arrival date onwards. However, the seller may find it suboptimal to participate at date 1, while the buyer may find it suboptimal to participate at the arrival date τ_B . We now explain how to ensure participation constraints bind at each date for the lowest realization of the buyer's value,²⁴ with the implication that the seller is willing to participate at date 1 for all realizations of his date-1 type $\theta_{S,1}$.

Let Ω^{TEAM} be the balanced team mechanism described above. Let $V_{B,\tau_B}^{\Omega^{TEAM}}(\theta_{B,\tau_B}; \theta_{S,\tau_B-1})$ be the expected payoff of a buyer who arrives at τ_B with value θ_{B,τ_B} when the seller has reported θ_{S,τ_B-1} in the previous period. We can modify the balanced team mechanism Ω^{TEAM} as follows. First require that the buyer, if he arrives at date 1, make a payment to the seller at date 1 equal to

$$V_{B,1}^{\Omega^{TEAM}}(\underline{\theta}_{B,1}) - \delta \mathbb{E} \left[V_{B,2}^{\Omega^{TEAM}}(\tilde{\theta}_{B,2}; \tilde{\theta}_{S,1}) \mid \tilde{\theta}_{B,1} = \underline{\theta}_{B,1} \right].$$

The buyer is then willing to participate at date 1 for all realizations of $\theta_{B,1}$, and earns the lowest payoff possible in a mechanism satisfying E, PIC and PPC, and which treats the buyer, if arriving at date 2 or later, precisely as in Ω^{TEAM} . We then say that the mechanism has been "modified" at date 1.

We define subsequent modifications, i.e. for $t \geq 2$, as follows. Suppose the mechanism has been modified up to date $t - 1$, and in particular that the mechanism satisfies E, PIC and PPC, while ensuring the buyer's rents if participating at date $t - 1$ or earlier are as small as possible, subject to E, PIC and PPC, and to the mechanism played by buyers participating from t onwards being specified by Ω^{TEAM} . Then require any buyer who participates at date $s \leq t$ to pay the seller an amount, conditional on $\theta_{S,1}^{s-1}$, equal to

$$\begin{aligned} & \Gamma_{s,t}(\theta_{S,1}^{s-1}) \\ \equiv & \delta^{t-s} \mathbb{E} \left[V_{B,t}^{\Omega^{TEAM}}(\underline{\theta}_{B,t}; \tilde{\theta}_{S,t-1}) - \delta V_{B,t+1}^{\Omega^{TEAM}}(\tilde{\theta}_{B,t+1}; \tilde{\theta}_{S,t}) \mid \tilde{\theta}_{B,t} = \underline{\theta}_{B,t}, \tilde{\theta}_{S,s-1} = \theta_{S,s-1} \right] \end{aligned}$$

on the participation date. Note that this ensures the buyer is willing to participate

²⁴That the buyer gains least from participating immediately (rather than waiting until the next period) when his value is at its lowest (i.e., $\underline{\theta}_{B,\tau_B}$ for each participation date τ_B) can be seen from considering the repetition of (static) VCG mechanisms and using the payoff equivalence property of Assumption 1.

for all arrival times up to and including t (and report truthfully thereafter), while ensuring that his expected payoff at each participation date $s \leq t$ is as small as possible in an efficient mechanism where buyers participating at date $t + 1$ or later receive the same treatment as under Ω^{TEAM} . Note that the latter holds because participation constraints bind at $\underline{\theta}_{B,s}$ for each date $s \leq t$ and for each possible history of seller reports $\theta_{S,1}^{s-1}$.

Note then that the aforementioned payments affect the seller's incentives to report truthfully at dates $t - 1$ and earlier. To resolve this, for each date s from 2 up to t , have the buyer, if participating at $s - k$ for $k \in \{1, 2, \dots, s - 1\}$, pay on participation the amount

$$\frac{\delta^k g_B(s)}{g_B(s-k)} \left(\mathbb{E} \left[\Gamma_{s,t} \left(\tilde{\theta}_{S,1}^{s-1} \right) \mid \tilde{\theta}_{S,1}^{s-k-1} = \theta_{S,1}^{s-k-1} \right] - \mathbb{E} \left[\Gamma_{s,t} \left(\tilde{\theta}_{S,1}^{s-1} \right) \mid \tilde{\theta}_{S,1}^{s-k} = \theta_{S,1}^{s-k} \right] \right). \quad (21)$$

Note that these payments are constructed to precisely "undo" the effect on the seller's incentive to misreport of the payment $\Gamma_{s,t}(\theta_{S,1}^{s-1})$ made by the buyer when participating at $s \leq t$.

By the law of iterated expectations, the expectation of the payment given in (21), conditional on $\theta_{S,1}^{s-k-1}$ for the buyer in case arriving at $s - k$, is equal to zero. Hence, the additional payments in (21) do not affect the buyer's expected payoff given that the seller truthfully reports his costs. Moreover, the payments in (21) are independent of the buyer's reported values, and do not affect incentives to report these truthfully. It follows that the modified mechanism is E, PIC and PPC, and ensures the buyer earns the smallest rents possible in an efficient mechanism which subjects any participant at $t + 1$ or later to the original mechanism Ω^{TEAM} . Proceeding this way, we inductively define a public mechanism satisfying E, PIC and PPC, and such that the buyer's participation constraints are binding at each date t for the lowest buyer value $\underline{\theta}_{B,t}$. It is then easy to see that the buyer's expected payoffs are as small as possible in a public mechanism satisfying E, PIC and PPC.²⁵

Finally, suppose we can show that the expected present value of the seller's payoff

²⁵To see this, note that the above defines a mechanism which, when blind, satisfies E, BIC and BPC, and is such that buyer participation constraints bind at $\underline{\theta}_{B,t}$ for any possible participation date $t \in \mathbb{N}$. Using that the buyer's expected payoffs from participating remain bounded across participation dates t , and the payoff equivalence property of Definition 1, the above blind mechanism must imply the minimal buyer rents among all blind mechanisms satisfying E, BIC and BPC (in particular, minimizing expected rents at each participation date t). But buyer expected rents cannot be lower in a public mechanism satisfying E, PIC and PPC, which implies the result.

at date 1, conditional on each $\theta_{S,1} \in \Theta_{S,1}$, is well-defined (so that the expected present value of the buyer and seller's combined payoffs, conditional on θ_1 , equal the expected present value of surplus from efficient trade). Then, using payoff equivalence to a sequence of VCG mechanisms (with fixed participation fees paid by the buyer to ensure participation constraints always bind at the lowest value), the seller's date-1 expected payoff when his initial cost is $\bar{\theta}_{S,1}$ must equal $\bar{U}_1 \geq 0$.²⁶ Assuming the seller may only participate at date 1 (so that his payoff from not participating at date 1 is zero), the seller's participation constraint is satisfied.

It therefore remains to check that the seller's payoff is well-defined (and hence NPS is satisfied), which is true if the expected present value of payments to the seller is well-defined. First, note that a buyer who participates at date a makes payments associated with the team mechanism which are uniformly bounded, and hence have a date- a present value that is uniformly bounded across all realizations of uncertainty. The expected present value of payments to the seller in the team mechanism are therefore well-defined. If the buyer participates at date a , then he also makes a date- a payment given the realization of seller costs $\theta_{S,1}^a$ equal to the amount

$$\sum_{t=a}^{\infty} \delta^{t-a} \left\{ \begin{array}{l} \mathbb{E} \left[V_{B,t}^{\Omega^{TEAM}} \left(\underline{\theta}_{B,t}; \tilde{\theta}_{S,t-1} \right) - \delta V_{B,t+1}^{\Omega^{TEAM}} \left(\tilde{\theta}_{B,t+1}; \tilde{\theta}_{S,t} \right) \mid \tilde{\theta}_{B,t} = \underline{\theta}_{B,t}, \tilde{\theta}_{S,a-1} = \theta_{S,a-1} \right] \\ + \frac{G_B(t) - G_B(a)}{g_B(a)} \left(\begin{array}{l} \mathbb{E} \left[\begin{array}{l} V_{B,t}^{\Omega^{TEAM}} \left(\underline{\theta}_{B,t}; \tilde{\theta}_{S,t-1} \right) \\ - \delta V_{B,t+1}^{\Omega^{TEAM}} \left(\tilde{\theta}_{B,t+1}; \tilde{\theta}_{S,t} \right) \end{array} \mid \begin{array}{l} \tilde{\theta}_{B,t} = \underline{\theta}_{B,t}, \\ \tilde{\theta}_{S,a-1} = \theta_{S,a-1} \end{array} \right] \\ - \mathbb{E} \left[\begin{array}{l} V_{B,t}^{\Omega^{TEAM}} \left(\underline{\theta}_{B,t}; \tilde{\theta}_{S,t-1} \right) \\ - \delta V_{B,t+1}^{\Omega^{TEAM}} \left(\tilde{\theta}_{B,t+1}; \tilde{\theta}_{S,t} \right) \end{array} \mid \begin{array}{l} \tilde{\theta}_{B,t} = \underline{\theta}_{B,t}, \\ \tilde{\theta}_{S,a} = \theta_{S,a} \end{array} \right] \end{array} \right) \end{array} \right\}.$$

We can then note that, because payoff types are uniformly bounded, and by definition of the balanced team mechanism,

$$V_{B,t}^{\Omega^{TEAM}} \left(\underline{\theta}_{B,t}; \theta_{S,t-1} \right) - \delta \mathbb{E} \left[V_{B,t+1}^{\Omega^{TEAM}} \left(\tilde{\theta}_{B,t+1}; \tilde{\theta}_{S,t} \right) \mid \tilde{\theta}_{B,t} = \underline{\theta}_{B,t}, \tilde{\theta}_{S,t-1} = \theta_{S,t-1} \right]$$

is uniformly bounded across all t and $\theta_{S,t-1} \in \Theta_{S,t-1}$. Hence, the expected present value of the absolute value of all buyer payments (conditional on each $\theta_{S,1} \in \Theta_{S,1}$) is finite, and the expected present value of buyer payments is well-defined. ■

²⁶ Given Assumption 1, payoff equivalence must be applied by considering these mechanisms played blind. That is, we can consider payoff equivalence between the constructed mechanism, which is BIC, E, and yields the lowest rents to buyers, and the aforementioned sequence of VCG mechanisms, which has the same properties.

Proof of Example 1. Follows from arguments in the main text. ■

Proof of Example 2. Consider the mechanism in which (i) no trade occurs and no payments are made until one period or more after both agents have participated, and (ii) after both agents have participated for at least one period, an AGV mechanism is played, with the buyer receiving an additional amount from the seller on each such date which ensures his ex-ante expected payoff from each AGV mechanism is zero. After an agent has participated, he is bound to participate forever after. The above mechanism is clearly E and BIC in the blind environment or PIC in the public environment. While the buyer's ex-ante expected payoff from each AGV mechanism is zero, the seller's expected payoff is zero conditional on the highest cost $\bar{\theta}$ (i.e., the seller's ex-ante expected payoff from each AGV mechanism is positive). It follows that the buyer expects zero from participating in the mechanism (either if the mechanism is blind, or if it is public then for any history of seller reports). The seller strictly prefers to participate at the first opportunity (either if the mechanism is blind, or if it is public then for any history of buyer reports). ■

Proof of Example 3. The expression for agent rents in (16) and (17) follow from the same arguments as in Garrett (2016), and are explained in the main text. The "relaxed program" for choosing the constrained-efficient allocation then consists of maximizing (15) subject to (18). The Lagrangian objective for this problem is

$$\mathbb{E} \left[\sum_{t=\max\{\tilde{\tau}_S, \tilde{\tau}_B\}}^{\infty} \delta^{t-1} (\tilde{\theta}_{B,t} - \tilde{\theta}_{S,t}) x_t (\tilde{\theta}_{B,\tau_B}^t, \tilde{\theta}_{S,\tau_S}^t) \right] + \lambda \left(\mathbb{E} \left[\sum_{t=\max\{\tilde{\tau}_B, \tilde{\tau}_S\}}^{\infty} \delta^{t-1} (\tilde{\theta}_{B,t} - \tilde{\theta}_{S,t}) x_t (\tilde{\theta}_{B,\tilde{\tau}_B}^t, \tilde{\theta}_{S,\tilde{\tau}_S}^t) \right] - \mathbb{E} \left[\begin{aligned} &\delta^{\tilde{\tau}_B-1} \frac{G_B(\tilde{\tau}_B)}{g_B(\tilde{\tau}_B)} \frac{1-F_B(\tilde{\theta}_{B,\tilde{\tau}_B})}{f_B(\tilde{\theta}_{B,\tilde{\tau}_B})} x_{\tilde{\tau}_B} (\tilde{\theta}_{B,\tilde{\tau}_B}, \tilde{\theta}_{S,\tilde{\tau}_S}^{\tilde{\tau}_B}) \\ &+ \delta^{\tilde{\tau}_S-1} \frac{G_S(\tilde{\tau}_S)}{g_S(\tilde{\tau}_S)} \frac{F_S(\tilde{\theta}_{S,\tilde{\tau}_S})}{f_S(\tilde{\theta}_{S,\tilde{\tau}_S})} x_{\tilde{\tau}_S} (\tilde{\theta}_{B,\tilde{\tau}_B}^{\tilde{\tau}_S}, \tilde{\theta}_{S,\tilde{\tau}_S}) \end{aligned} \right] \right) \quad (22)$$

for a multiplier λ . Permitting the allocation x_t to take values in $[0, 1]$ for all t , by linearity, the optimization of (22) has a bang-bang solution for each $\lambda \geq 0$ given by (19). Given that (18) varies continuously in the allocation, and given that the constraint (18) fails for the efficient allocation rule (by Proposition 2, since the payoff-type supports overlap), there exists a unique value for the multiplier, $\lambda^* > 0$, such that (18) holds with equality.

The result in the example then follows if a public mechanism implementing $(x_t^*)_{t \geq 1}$ can be found satisfying PIC, PPC and EBS. The possibility to find a PIC mechanism follows because the allocation $(x_t^*)_{t \geq 1}$ is monotone in agents' payoff types, which follows from the monotonicity of $\frac{1-F_B(\cdot)}{f_B(\cdot)}$ and $\frac{F_S(\cdot)}{f_S(\cdot)}$. Transfers can then be adjusted by fixed (type-independent) payments such that, if agent i participates for the first time at date τ_i with i 's "worst" payoff type ($\underline{\theta}_B$ for the buyer or $\bar{\theta}_S$ for the seller), he expects the same payoff as from delaying participation by exactly one period (where this expectation is calculated conditional on agent $-i$'s reports up to $\tau_i - 1$). In particular, this expected payoff is set to equal

$$\sum_{t=\tau_B+1} \delta^{t-\tau_B} \mathbb{E} \left[\frac{1 - F_B(\tilde{\theta}_{B,t})}{f_B(\tilde{\theta}_{B,t})} x_t^* \left(\tilde{\theta}_{B,t}, \tilde{\theta}_{S,\tilde{\tau}_S}^t \right) \mid \tilde{\theta}_{S,\tilde{\tau}_S}^{\tau_B-1} = \theta_{S,\tilde{\tau}_S}^{\tau_B-1}, \tilde{\theta}_{B,t} \neq \emptyset \right]$$

for the buyer participating at τ_B (when his value is $\underline{\theta}_B$ and the seller has made reports $\theta_{S,\tilde{\tau}_S}^{\tau_B-1}$ up to $\tau_B - 1$), and to

$$\sum_{t=\tau_S+1} \delta^{t-\tau_S} \mathbb{E} \left[\frac{F_S(\tilde{\theta}_{S,t})}{f_S(\tilde{\theta}_{S,t})} x_t^* \left(\tilde{\theta}_{B,\tilde{\tau}_B}^t, \tilde{\theta}_{S,t} \right) \mid \tilde{\theta}_{B,\tilde{\tau}_B}^{\tau_S-1} = \theta_{B,\tilde{\tau}_B}^{\tau_S-1}, \tilde{\theta}_{S,t} \neq \emptyset \right]$$

for the seller participating at τ_S (when his cost is $\bar{\theta}_S$ and the buyer has made reports $\theta_{B,\tilde{\tau}_B}^{\tau_S-1}$ up to $\tau_S - 1$). The adjusted mechanism is then PIC and PPC and leads to the smallest expected rents for the agents among mechanisms implementing $(x_t^*)_{t \geq 1}$ (whether among blind mechanisms satisfying BIC and BPC or public mechanisms satisfying PIC and PPC). Expected agent rents are then given by (16) and (17), and the constraint (18) is satisfied with equality, meaning that constraint EBS is satisfied.

■