HARVARD | BUSINESS | SCHOOL

Bundling Incentives in (Many-to-Many) Matching with Contracts

Jonathan Ma
Scott Duke Kominers

Working Paper 19-011

Bundling Incentives in (Many-to-Many) Matching with Contracts

Jonathan Ma
Harvard University

Scott Duke Kominers
Harvard Business School

Working Paper 19-011

Bundling Incentives in (Many-to-Many) Matching with Contracts*

Jonathan Ma
Department of Applied Mathematics
Harvard University

Scott Duke Kominers
Harvard Business School and Department of Economics
Harvard University

August 11, 2018

Abstract

In many-to-many matching with contracts, the way in which contracts are specified can affect the set of stable equilibrium outcomes. Consequently, agents may be incentivized to modify the set of contracts upfront. We consider one simple way in which agents may do so: *unilateral bundling*, in which a single agent links multiple contracts with the same counterparty together. We show that essentially no stable matching mechanism eliminates incentives for unilateral bundling. Moreover, we find that unilateral bundling can sometimes lead to Pareto improvement—and other times produces market power that makes one agent better off at the expense of others.

JEL Classification: C62; C78; D44; D47

Keywords: Matching with contracts; Contract design; Bundling-proofness; Substitutability

*The authors thank Jiafeng Chen, John William Hatfield, Ravi Jagadeesan, Jeff Miron, and Alexandru Nichifor for helpful comments. Kominers gratefully acknowledges the support of the National Science Foundation (grant SES-1459912), as well as the Ng Fund and the Mathematics in Economics Research Fund of the Harvard Center of Mathematical Sciences and Applications.

1 Introduction

The *matching with contracts* model (Crawford and Knoer (1981); Kelso and Crawford (1982); Hatfield and Milgrom (2005)) unifies three disparate market design frameworks: the college admissions/marriage model of Gale and Shapley (1962), the labor market/assignment matching model of Shapley and Shubik (1971), Crawford and Knoer (1981), and Kelso and Crawford (1982), and the ascending package auction models studied by Ausubel and Milgrom (2002). In recent years, matching with contracts has been used to analyze and design markets ranging from the US Military cadet–branch matches (Sönmez and Switzer (2013); Sönmez (2013)) to the Israeli Psychology Masters Match (Hassidim et al., 2017). Applying the matching with contracts model requires a degree of subtlety, however, as the way in which contracts are specified affects the analysis (Hatfield and Kominers (2015)), and can even change the set of equilibria (Hatfield and Kominers (2017)).

Hatfield and Kominers (2017) introduced a framework for analyzing *contract design* in matching settings, and showed how different ways of grouping the same set of *contractual primitives* into contracts could affect the set of stable matching outcomes. Here, we build upon the work of Hatfield and Kominers (2017), exploring the strategic incentives that might result from different choices of contract language. To develop this idea, we consider a form of unilateral *bundling* that has been studied in the context of exchange markets with indivisible goods.

Klaus et al. (2006) demonstrated that agents in exchange markets may have incentives to bundle indivisible goods together in order to constrain the set of possible allocations in ways that generate market power. We find a similar result in our context: If given the opportunity, agents may prefer to modify the contract language (prior to matching) by bundling a set of contracts into a single contract—and thus constraining the set of stable outcomes. Sometimes, unilateral bundling of contracts can lead to Pareto improvement; other times, it produces market power that makes one agent better off at the expense of others.

The remainder of this paper is organized as follows. Section 2 reviews the model of (many-to-many) matching with contracts, largely following the contract language framework of Hatfield and Kominers (2017). In Section 3, we extend the matching process by adding an *ex ante* stage in which agents are able to bundle contracts unilaterally. In Section 4, we build upon an example from Hatfield and Kominers (2017) to demonstrate that in a market with only two agents, it is possible that both agents would be incentivized to bundle a set of contracts *ex ante*. We then show with a second example that even in the two-agent case, we may construct preferences over contractual primitives such that one agent is incentivized to bundle a set of contracts, while the other agent is not. By extension, our findings imply that no stable matching mechanism eliminates unilateral bundling incentives; that is, no stable matching mechanism is *bundling-proof*.

2 Many-to-Many Matching with Contracts

2.1 Basic Framework

We begin with the many-to-many matching with contracts model of Hatfield and Kominers (2017), using the same notation as much as possible (see also Klaus and Walzl (2009)).

There is a set D of *doctors* and a set H of *hospitals*; the set of *agents* is $F \equiv D \cup H$. A set of contracts between doctor–hospital pairs is denoted by X. Each contract $x \in X$ is then associated with a doctor $x_D \in D$ and a hospital $x_H \in H$; we define $x_F \equiv \{x_D, x_H\}$ to be the set of agents associated with contract $x \in X$.[1] Extending these notations, for any set of contracts $Y \in X$ we write

$$Y_D \equiv \bigcup_{y \in Y} \{y_D\}, \quad Y_H \equiv \bigcup_{y \in Y} \{y_H\}, \quad \text{and} \quad Y_F \equiv Y_D \cup Y_H,$$

respectively, for the sets of doctors, hospitals, and agents associated with contracts in Y.

For any $Y \subseteq X$ and $f \in F$, we write $Y_f \equiv \{y \in Y : f \in y_F\}$. We assume throughout

[1]For concreteness, we may treat X as a subset of $D \times H \times T$, where T is a set of possible contractual terms. In the sequel, we construct X explicitly from sets of contractual primitives.

that each agent $f \in F$ has a strict preference relation P_f^X (with associated weak relation R_f^X) over subsets of X_f. We assume that the preferences of each agent $f \in F$ are *non-unitary* (Kominers (2012)), in the sense that it is in principle possible for a given doctor–hospital pair to sign multiple contracts at the same time.

For any agent f and *offer set* $Y \subseteq X$, we can define a choice function

$$C_f(Y) \equiv \max_{P_f^X}\{Z \subseteq X : Z \subseteq Y_f\}$$

that specifies the set of contracts f chooses from Y.[2]

2.1.1 Stable Outcomes

An *outcome* is a set of contracts $Y \subseteq X$. Preferences extend naturally to outcomes: For $Y \subseteq X$ and $Z \subseteq X$, we say that $Y P_f^X Z$ if and only if $Y_f P_f^X Z_f$.

We again adopt the terminology and notation of Hatfield and Kominers (2017) characterizing outcomes. An outcome Y is *individually rational for* $f \in F$ if $C_f(Y) = Y_f$—that is, if f does not want to unilaterally abrogate any of his or her assigned contracts. An outcome Y is *unblocked* if there does not exist a nonempty *blocking set* $Z \subseteq [X \setminus Y]$ such that $Z_f \subseteq C_f(Y \cup Z)$ for all $f \in Z_F$—that is, if there is no set of agents who can profitably recontract among themselves (possibly while keeping some of their other contracts from Y). An outcome $Y \subseteq X$ is *stable* if it is

- individually rational for all $f \in F$ and
- unblocked.[3]

[2]The notation $\max_{P_f^X}$ indicates that the maximization is taken with respect to the preferences P_f^X of agent f.

[3]Our stability concept allows blocking deviations in which *sets* of agents recontract. Two other common stability concepts are (a) *many-to-one stability* and (b) *pairwise stability*, which respectively add to our definition the requirements (a) that blocking sets Z are associated to a single hospital, i.e., $|Z_H| = 1$, and (b) that blocking sets Z consist of at most one contract, i.e., $|Z| = 1$. Under substitutable preferences (defined in the next section), our stability concept, many-to-one stability, and pairwise stability coincide (Hatfield and Kominers (2017)).

2.1.2 Substitutable Preferences

Following Hatfield and Milgrom (2005) (and Kelso and Crawford (1982), Roth (1984), and other antecedents; see also Hatfield et al. (2018)), we say that the preferences of f are *substitutable* if, for all $x, z \in X$ and $Y \subseteq X$,

$$z \notin C_f(Y \cup \{z\}) \implies z \notin C_f(\{x\} \cup Y \cup \{z\}).$$

Substitutability means that there are no two contracts x and z that are *complements*, in the sense that getting access to x makes z more desirable.

2.2 Contract Language

Again following Hatfield and Kominers (2017), we model the agents' preferences P_f^X over contracts as derived from preferences over underlying primitives that are bundled into contracts. For each doctor–hospital pair $(d, h) \in D \times H$, we assume that there is a set of *contractual primitives* $\pi(d, h)$ that defines all possible contractual relationships between d and h. Each contract between d and h is a subset of $\pi(d, h)$. With this structure, the set of contracts between d and h,

$$X_{(d,h)} \equiv (X_d)_h = (X_h)_d,$$

is a subset of $\mathcal{P}(\pi(d, h)) \setminus \{0\}$, where the notation $\mathcal{P}$ denotes the power set. We define a *contract language* $X_{(d,h)}$ *for* $(d, h) \in (D, H)$ as a set of contracts between d and h; therefore, $X_{(d,h)}$ is a subset of $\mathcal{P}(\pi(d, h))\setminus\{0\}$. A given contract set $X = \cup_{(d,h)\in D\times H} X_{(d,h)}$ is called a *contract language.*

2.2.1 Primitive Outcomes

A *primitive outcome* is a collection of primitives

$$\Lambda \subseteq \bigcup_{(d,h)\in D\times H} \pi(d,h).$$

We say that a primitive outcome Λ is *expressible* in the contract language X if we have $\Lambda = \cup_{y\in Y} y$ for some $Y \subseteq X$; we say in that case that Y *expresses* Λ.

2.2.2 Preferences over Primitives

We write $\Pi_d \equiv \cup_{h\in H}\pi(d,h)$ for the set of primitives associated with doctor $d \in D$; analogously, we write $\Pi_h \equiv \cup_{d\in D}\pi(d,h)$ for the set of primitives associated with hospital $h \in H$. We consider each agent $f \in F$ to have a strict preference relation P_f over the set $\mathcal{P}(\Pi_f)$ of sets of primitives involving f. This preference relation P_f induces the agent's preference relation P_f^X over sets of contracts in X_f, via the relation

$$Y P_f^X Z \iff \left[\bigcup_{y\in Y_f} y\right] P_f \left[\bigcup_{z\in Z_f} z\right];^{4}$$

that is, we require that each agent f has a preference relation over sets of contracts, P_f^X, consistent with his or her preference relation over sets of primitives, P_f.[5]

3 Incentives to Bundle Contracts

Our framework thus far—as well as the prior work on contract language design—has treated the contract language as fixed by the market/mechanism designer. However, Hatfield and

[4]To complete P_f^X, we assume that $Y R_f^X Z \iff \left[\cup_{y\in Y_f} y\right] = \left[\cup_{z\in Z_f} z\right]$.

[5]Notably, the induced preference relation over sets of contracts may not be strict because two sets of contracts may be *primitive-equivalent*, in the sense that they express the same primitive outcome. When defining the choice function C_f, we break ties between primitive-equivalent sets of contracts arbitrarily in a way that makes C_f consistent with P_f.

Kominers (2017) showed that *bundling* individual contracts together can affect the set of stable outcomes. In other words, if we choose some $x, z \in X_{(d,h)}$ and replace the contract language X with $\hat{X} = [X \setminus \{x, z\}] \cup \{x \cup z\}$, then the set of stable outcomes may change. Because the set of stable outcomes can be affected by bundling, agents may be incentivized to bundle contracts in order to eliminate stable outcomes they view as unfavorable. In real markets, agents often have leeway to aggregate multiple transactions with the same transaction partner; hence, it makes sense to consider the possibility of bundling as a strategic action.

3.1 The Modified Matching Process

We suppose that agents have some ability to affect the contract language before the matching mechanism is run. Specifically, we assume that each agent f can *unilaterally bundle* some of his or her contracts in X, that is, f may choose a set of contracts $Y^f \subseteq X$, removing all $y \in Y^f$ from the contract language,[6] and add to the contract language a single contract $\cup_{y \in Y} y$ representing the "bundle" containing all the primitives associated to contracts in Y.[7]

Formally, our modified matching process proceeds as follows:

0. Given a mechanism $\mathcal{M}$ and contract language X,

1. all agents $f \in F$ (simultaneously) choose (possibly empty) sets of contracts $Y^f \subseteq X$ to unilaterally bundle, resulting in a new contract set

$$\hat{X} = \left[X \setminus \left[\cup_{f \in F} Y^f\right]\right] \bigcup \left[\cup_{f \in F} \left[\cup_{y \in Y^f} y\right]\right];$$

after that,

[6] We model bundling as removing the union of all the individual contracts combined in the bundle—$\cup_{f \in F} Y^f$—from the contract language, as we believe this best models real-world contracting scenarios. All of our results continue to hold in an alternate model of bundling in which the "bundle" contract is added to the contract set while leaving the bundle's component contracts available; see Appendix A.

[7] We consider only the possibility that each f may introduce a single bundle; since our results are negative, they extend to a more general concept of bundling in which each agent is allowed to create multiple distinct bundles unilaterally.

2. each agent f submits his or her preferences $P_f^{\hat{X}}$ over contracts in $\hat{X}$ to the mechanism $\mathcal{M}$, and $\mathcal{M}$ selects the outcome $\mathcal{M}(P^{\hat{X}}) \subseteq \hat{X}$.

Note that it is possible that $Y^i \cap Y^j \neq \emptyset$ for some agents $i \neq j$: that is, two agents may select the same contracts to bundle. Given agents i and j as well as contracts $x^\alpha, x^\beta, x^\gamma \in Y^i \cap Y^j$, for example, agent i may attempt to bundle contracts x^α and x^β while agent j may attempt to bundle contracts x^β and x^γ. For the remainder of this paper, we dictate that when $Y^i \cap Y^j \neq \emptyset$, both $\cup_{y \in Y^i} y$ and $\cup_{y \in Y^j} y$ are added to the contract language via setwise union.

We say that a mechanism is *stable* if it always yields a stable outcome (for any contract language, and any input preferences). We say that a mechanism is *bundling-proof* if there is no profile of preferences P (over primitives) and contract language X such that some agent f can obtain a more-preferred (primitive) outcome by unilaterally bundling some or all of his of her contracts in X.

4 Main Results

4.1 Pareto Improvement via Bundling

First, we build upon a two-agent example of Hatfield and Kominers (2017), in which both agents would be incentivized to bundle contracts, as doing so results in a Pareto improvement over the stable outcome under the original unbundled contract language.[8]

Example 1. Consider a doctor d and hospital h. Suppose that both d and h strictly disprefer all sets of contractual primitives to the empty set, except for the primitives $\alpha, \beta \in \pi(d, h)$.

[8]Rostek and Yoder (2017) introduced a closely related example in a very different context—specifically, comparison of the set of stable outcomes with the set of *setwise stable* outcomes (Echenique and Oviedo (2006); Klaus and Walzl (2009)). Like in our example, in Rostek and Yoder's (2017) Example 2, agents disagree about which contracts to keep from a set that is not individually rational; this means that the set of stable outcomes is Pareto inferior to the outcome that would arise if the agents could commit to keep the same set of contracts (as they can under bundling, or under the setwise stability solution concept).

Moreover, suppose that d and h have preferences over α and β as follows:

$$P_d : \{\beta\} \succ \{\alpha, \beta\} \succ \varnothing \succ \{\alpha\}$$
$$P_h : \{\alpha\} \succ \{\alpha, \beta\} \succ \varnothing \succ \{\beta\}.$$

For concreteness, we can imagine, for example, that β represents "work" that the doctor does for the hospital and α represents the hospital's compensation package for the doctor. Of course, in his or her heart-of-hearts, the doctor would most like to be paid for nothing; the hospital would most like the doctor to work for free.

Abusing notation slightly, we follow Hatfield and Kominers (2017) in writing x^Γ for the contract $\{\Gamma\}$. If the contract language X includes the contracts x^α and x^β but does not contain $x^{\alpha,\beta}$, then the preferences over contracts induced by d's and h's preferences over primitives are given by

$$P_d^X : \{x^\beta\} \succ \{x^\alpha, x^\beta\} \succ \varnothing \succ \{x^\alpha\}$$
$$P_h^X : \{x^\alpha\} \succ \{x^\alpha, x^\beta\} \succ \varnothing \succ \{x^\beta\}.$$

Note that something seemingly pathological happens: Under the preferences P^X, the unique stable outcome is $\varnothing$, as $\{x^\alpha, x^\beta\}$ is not individually rational for either agent.

However, if x^α and x^β are bundled into a single contract to yield the new contract language

$$\hat{X} = [X \setminus \{x^\alpha, x^\beta\}] \cup \{x^\alpha \cup x^\beta\} = [X \setminus \{x^\alpha, x^\beta\}] \cup \{x^{\alpha,\beta}\},$$

then the agents' preferences over outcomes become

$$P_d^{\hat{X}} : \{x^{\alpha,\beta}\} \succ \varnothing$$
$$P_h^{\hat{X}} : \{x^{\alpha,\beta}\} \succ \varnothing,$$

under which the unique stable outcome is $\{x^{\alpha,\beta}\}$. As $\{\alpha,\beta\}P_d\varnothing$ and $\{\alpha,\beta\}P_h\varnothing$, bundling x^α and x^β results in a (strict) Pareto improvement (under the primitive preferences of the agents), for any stable matching mechanism. Thus both d and h are incentivized to unilaterally bundle x^α and x^β.

While the unbundled preferences over contracts P_d^X and P_h^X in Example 1 are substitutable, substitutability of preferences over contracts is not a necessary condition for Pareto improvement via bundling, as we see in the following example.

Example 2. Maintaining the setting of Example 1, suppose that the agents' preferences over primitives are now as follows:

$$P_d : \{\alpha\} \succ \{\alpha,\beta\} \succ \{\beta\} \succ \varnothing$$
$$P_h : \{\alpha,\beta\} \succ \varnothing \succ \{\alpha\} \succ \{\beta\}.$$

Thus the agents' preferences over contracts in the unbundled contract language X are

$$P_d^X : \{x^\alpha\} \succ \{x^\alpha, x^\beta\} \succ \{x^\beta\} \succ \varnothing$$
$$P_h^X : \{x^\alpha, x^\beta\} \succ \varnothing \succ \{x^\alpha\} \succ \{x^\beta\};$$

the preferences of h are not substitutable, and as in Example 1, $\varnothing$ is the only stable outcome under X. However, under the bundled contract language $\hat{X}$, the agents' preferences over outcomes become

$$P_d^{\hat{X}} : \{x^{\alpha,\beta}\} \succ \varnothing$$
$$P_h^{\hat{X}} : \{x^{\alpha,\beta}\} \succ \varnothing,$$

so that $\{x^{\alpha,\beta}\}$ is the unique stable outcome under $\hat{X}$. As both agents strictly prefer $\{\alpha,\beta\}$ to $\varnothing$, bundling x^α and x^β again results in a (strict) Pareto improvement—under the agents' primitive preferences—under any stable matching mechanism. We conclude that even when agents' preferences over contracts are not substitutable, agents may still be incentivized to engage in unilateral bundling.

4.2 Conflicting Incentives to Bundle Contracts

Now, we demonstrate that as long as there exist two contractual primitives associated with two agents whose preferences are nonunitary, it is always possible to construct preferences over primitives such that one agent has an incentive to bundle unilaterally, while the other agent would prefer not to do so.

Example 3. Again using the setting of Example 1, we suppose that the agents' preferences over primitives are as follows:

$$P_d : \{\alpha\} \succ \{\alpha, \beta\} \succ \{\beta\} \succ \varnothing$$
$$P_h : \{\alpha, \beta\} \succ \{\alpha\} \succ \{\beta\} \succ \varnothing.$$

Then in a contract language X in which α and β are unbundled, the agents' preferences take the form

$$P_d^X : \{x^\alpha\} \succ \{x^\alpha, x^\beta\} \succ \{x^\beta\} \succ \varnothing$$
$$P_h^X : \{x^\alpha, x^\beta\} \succ \{x^\alpha\} \succ \{x^\beta\} \succ \varnothing;$$

the unique stable outcome under X is $\{x^\alpha\}$. However, if α and β are bundled into a single contract to yield the new contract language $\hat{X}$, then the agents' preferences over outcomes again become

$$P_d^{\hat{X}} : \{x^{\alpha,\beta}\} \succ \varnothing$$
$$P_h^{\hat{X}} : \{x^{\alpha,\beta}\} \succ \varnothing,$$

with the unique stable outcome $\{x^{\alpha,\beta}\}$, which h prefers to $\{x^{\alpha}\}$, but d likes less than $\{x^{\alpha}\}$. Thus, h is incentivized to unilaterally bundle x^{α} and x^{β}, while d is not.

Like in Example 1, the preferences in the Example 3 are substitutable, but as with the case of Pareto improvements through bundling, we can show with an example that agents may have conflicting incentives to bundle even in the absence of substitutability.

Example 4. We once more use the setting of Example 1, but now we specify the agents' preferences over primitives as follows:

$$P_d : \{\alpha\} \succ \{\alpha, \beta\} \succ \{\beta\} \succ \varnothing$$
$$P_h : \{\alpha, \beta\} \succ \{\alpha\} \succ \varnothing \succ \{\beta\}.$$

In the case in which α and β are unbundled, the preferences of each agent over the possible outcomes take the form:

$$P_d^X : \{x^{\alpha}\} \succ \{x^{\alpha}, x^{\beta}\} \succ \{x^{\beta}\} \succ \varnothing$$
$$P_h^X : \{x^{\alpha}, x^{\beta}\} \succ \{x^{\alpha}\} \succ \varnothing \succ \{x^{\beta}\};$$

these preferences over contracts are not substitutable. The unique stable outcome under X is $\{x^{\alpha}\}$. However, just as in Example 3, when α and β are bundled into a single contract under the new contract language $\hat{X}$, the only stable outcome is $\{x^{\alpha,\beta}\}$. Thus, once more, h is incentivized to unilaterally bundle x^{α} and x^{β}, while d is not.

4.3 Generalization to All Matching Problems

Embedding our examples in a broader market context immediately implies that no stable matching mechanism is bundling-proof, in general.

Theorem 1. *Suppose that there exists a doctor d and hospital h who share at least two contractual primitives (i.e., $|\pi(d, h)| \geq 2$). Then, for any stable matching mechanism, there* *exists a profile of preferences over primitives P and a contract language X such that at least one of d and h would like to unilaterally bundle contracts in X prior to matching. That is, no stable mechanism is bundling-proof.*

A Appendix: An Alternate Model of Bundling

In Section 3.1, we took the contract language resulting from bundling to be

$$\hat{X} = \left[X \setminus \left[\cup_{f\in F} Y^f\right]\right] \bigcup \left[\cup_{f\in F} \left[\cup_{y\in Y^f} y\right]\right];$$

we remove the all the individual contracts that are combined into the bundle, $\cup_{f\in F} Y^f$, from the contract language. For instance, if a contract language for a doctor and hospital contains two contracts representing singleton primitives—the hospital paying the doctor and the doctor working for the hospital, as in Example 1—once we bundle those two primitives into a single contract, the two separate single-primitive contracts become unavailable.

Assuming that linking a set of contracts together renders the individual contracts unavailable as independent units is intuitive from a practical perspective, as that would reduce transaction costs. However, we could also imagine that bundling could occur without making constituent contracts unavailable; we show in this appendix that our results continue to hold under that alternate bundling model. We consider a mechanism that is identical to the one proposed in Section 3.1 except that the new contract language in Step 1 is redefined as

$$\hat{X} = X \bigcup \left[\cup_{f\in F} \left[\cup_{y\in Y^f} y\right]\right];$$

the contracts that are bundled together are not themselves removed from the contract language.

In Example 1, then the agents' preferences over contracts become

$$P_d^{\hat{X}} : \{x^\beta\} \succ \{x^\alpha, x^\beta\} \sim \{x^{\alpha,\beta}\} \succ \varnothing \succ \{x^\alpha\}$$
$$P_h^{\hat{X}} : \{x^\alpha\} \succ \{x^\alpha, x^\beta\} \sim \{x^{\alpha,\beta}\} \succ \varnothing \succ \{x^\beta\};$$

under these preferences the unique stable outcome is $\{x^{\alpha,\beta}\}$, just as in Example 1. It follows that both agents still have an incentive to bundle. If we similarly bundle the contracts of

Example 2 without removing the contracts chosen to be bundled, the agents' preferences over contracts become

$$P_d^{\hat{X}} : \{x^\alpha\} \succ \{x^\alpha, x^\beta\} \sim \{x^{\alpha,\beta}\} \succ \{x^\beta\} \succ \varnothing$$
$$P_h^{\hat{X}} : \{x^\alpha, x^\beta\} \sim \{x^{\alpha,\beta}\} \succ \varnothing \succ \{x^\alpha\} \succ \{x^\beta\};$$

these preferences again yield $\{x^{\alpha,\beta}\}$ as the unique stable outcome, just as in Example 2.

Moving to conflicting bundling incentives, suppose we bundle the contracts of Example 3 without removing the contracts chosen to be bundled. Then the preferences over the bundled contract language $\hat{X}$ are

$$P_d^{\hat{X}} : \{x^\alpha\} \succ \{x^\alpha, x^\beta\} \sim \{x^{\alpha,\beta}\} \succ \{x^\beta\} \succ \varnothing$$
$$P_h^{\hat{X}} : \{x^\alpha, x^\beta\} \sim \{x^{\alpha,\beta}\} \succ \{x^\alpha\} \succ \{x^\beta\} \succ \varnothing;$$

both $\{x^\alpha\}$ and $\{x^{\alpha,\beta}\}$ are stable under these preferences. Therefore, h is incentivized to unilaterally bundle x^α and x^β while d is not, as bundling x^α and x^β introduces the stable outcome $\{x^{\alpha,\beta}\}$, which is h's most preferred outcome—but d likes less than $\{x^\alpha\}$. Our conclusion from Example 3 therefore continues to hold. Likewise, in the setting of Example 4 we obtain the preferences

$$P_d^{\hat{X}} : \{x^\alpha\} \succ \{x^\alpha, x^\beta\} \sim \{x^{\alpha,\beta}\} \succ \{x^\beta\} \succ \varnothing$$
$$P_h^{\hat{X}} : \{x^\alpha, x^\beta\} \sim \{x^{\alpha,\beta}\} \succ \{x^\alpha\} \succ \varnothing \succ \{x^\beta\},$$

which yield the same conclusion.

As the logic of all of our examples holds under the alternate bundling model we consider in this appendix, Theorem 1—which extrapolates from those examples—does, as well.

References

Ausubel, L. M. and P. Milgrom (2002). Ascending auctions with package bidding. *Frontiers of Theoretical Economics 1*, 1–42.

Crawford, V. P. and E. M. Knoer (1981). Job matching with heterogeneous firms and workers. *Econometrica 49*(2), 437–50.

Echenique, F. and J. Oviedo (2006). A theory of stability in many-to-many matching markets. *Theoretical Economics 1*, 233–273.

Gale, D. and L. S. Shapley (1962). College admissions and the stability of marriage. *American Mathematical Monthly 69*, 9–15.

Hassidim, A., A. Romm, and R. I. Shorrer (2017). Redesigning the Israeli psychology master's match. *American Economic Review 107*(5), 205–09.

Hatfield, J. W. and S. D. Kominers (2015). Hidden substitutes. Harvard University Working Paper.

Hatfield, J. W. and S. D. Kominers (2017). Contract design and stability in many-to-many matching. *Games and Economic Behavior 101*, 78–97.

Hatfield, J. W., S. D. Kominers, A. Nichifor, M. Ostrovsky, and A. Westkamp (2018). Full subsitutability. Stanford University Working Paper.

Hatfield, J. W. and P. Milgrom (2005). Matching with contracts. *American Economic Review 95*, 913–935.

Kelso, A. S. and V. P. Crawford (1982). Job matching, coalition formation, and gross substitutes. *Econometrica 50*, 1483–1504.

Klaus, B., D. Dimitrov, and C.-J. Haake (2006). Bundling in exchange markets with indivisible goods. *Economics Letters 93*(1), 106–110.

Klaus, B. and M. Walzl (2009). Stable many-to-many matchings with contracts. *Journal of Mathematical Economics 45*(7-8), 422–434.

Kominers, S. D. (2012). On the correspondence of contracts to salaries in (many-to-many) matching. *Games and Economic Behavior 75*, 984–989.

Rostek, M. J. and N. Yoder (2017). Matching with multilateral contracts. University of Wisconsin Working Paper.

Roth, A. E. (1984). Stability and polarization of interests in job matching. *Econometrica 52*, 47–57.

Shapley, L. S. and M. Shubik (1971). The assignment game I: The core. *International Journal of Game Theory 1*, 111–130.

Sönmez, T. (2013). Bidding for army career specialties: Improving the ROTC branching mechanism. *Journal of Political Economy 121*, 186–219.

Sönmez, T. and T. B. Switzer (2013). Matching with (branch-of-choice) contracts at United States Military Academy. *Econometrica 81*, 451–488.