

*Negotiations in the Shadow of Outside Alternatives: An Estimation Strategy**

MARIUS RADEAN[†]
University of Essex

ABSTRACT

Many political and socio-economic binary outcomes are the result of multi-actor interaction: states joining a given international institution (e.g., military alliance, economic block) and not its rivals, people entering the workforce as an employee of a specific firm and not of its competitors, etc. Yet, most empirical studies analyze multilateral phenomena as the (joint) binary choice of either a single or, at most, two decision-makers. This is due in part to a lack of empirical estimators that can efficiently deal with multi-actor interaction. Analyzing multilateral processes as a set of either monadic or dyadic events, though, introduces bias, and has important consequences for the estimates and ultimately the inferences that one would draw. In this paper, I develop a new empirical estimator that is specifically designed to analyze multiparty interactions. Specifically, the model can accommodate the input of multiple actors into a unified, overarching decision making process. Results from a Monte Carlo analysis and an application to real data on alliance formation demonstrate the superior performance of the new estimator relative to the standard approach.

*Note: I thank Matt Golder, Daina Chiba, Garrett Glasgow, Kristian Skrede Gleditsch, James Honaker, Thomas Plümper, Christopher Reenock, and David Siegel for helpful comments on this project. The data and all computer code necessary to replicate the results and figures in this analysis will be made publicly available on my homepage on publication.

[†]Lecturer, Department of Government, University of Essex, Wivenhoe Park, Colchester, Essex, CO4 3SQ, United Kingdom, mradean@essex.ac.uk

1 Introduction

After the Ukrainian cabinet unanimously approved the E.U. Association Agreement in September 2013, the official signing of the agreement, scheduled two months later at the E.U. Vilnius summit, was supposed to be a simple formality. However, in November, the Ukrainian President Yanukovich refused to sign the agreement, a decision that triggered mass protests, and finally led to the current conflict in the eastern Ukraine. What derailed a cooperation that both parties found mutually beneficial only two months before? While examining the bilateral E.U.-Ukraine negotiations sheds some light into this sudden change of heart, it does not tell the complete story. Arguably, a concerned third party, Russia, interfered in, and ultimately altered, the E.U.-Ukraine relations. A combination of political pressure and economic side-payments, as well as the prospect of membership in the Eurasian Customs Union, led Ukraine closer to Russia (Englund & Lally 2013). Of course, closer ties with Russia meant weaker ties with the E.U. Thus, in order to understand the Vilnius debacle, politicians and scholars alike cannot focus on the E.U.-Ukraine interaction alone. They also need to take into account the actors' interaction with relevant third parties.

The decision that Ukraine faced in 2013 is by no means a singular event. Quite the contrary, situations in which a player simultaneously negotiates with more than one potential partner, and an agreement with one party precludes an agreement with the others, are ubiquitous. Governments that join a given international institution (e.g., military alliance, free trade area, etc.), implicitly decide against joining a rival organization. Employees who sign a contract with a specific employer often have alternative offers that they did not accept. Also, in most countries, people can marry only one person, politicians can be a member of only one political party, etc.

Even though multiparty negotiations are a common occurrence, analyzing such scenarios is hindered by a lack of empirical estimators that can efficiently deal with multi-actor interaction. In the absence of a well-established estimation strategy, empirical studies employ a wide range of alternative models (e.g., univariate logit/probit, (mixed) conditional and multinomial logit, bivariate probit, selection and partial observability models) to analyze binary outcomes that are the result of negotiations with outside alternatives. While in practice they are the most commonly used estimators, univariate and bivariate models share a crucial limitation: they can accommodate one or two actors, but not more (Greene 2003, Heckman 1979, McFadden 1974, Poirier 1980, Train 2007, Wooldridge 2003). Particularly in the context of multiparty nego-

tiations, this limitation has important consequences for the estimated parameters' consistency, and ultimately the inferences that one would draw. Specifically, results from Monte Carlo (MC) simulations indicate that ignoring the multilateral aspect in a multi-actor interaction scenario leads to biased results, and inaccurate or incorrect inferences.

This paper develops a new estimation procedure, called the Multilateral Mediated Interaction with Partial Observability (MMIwPO) model, for the analysis of binary decisions with competing and typically unobserved alternatives.¹ In this model, a player chooses whether to cooperate with any, or none, of several potential partners. For their part, each potential partner decides whether to reject or accept any requests for cooperation. If more than one suitable partner is willing to cooperate, the player chooses the partner that they prefer the best. Given that any potential partner can reject the player, the partner with which the player ends up cooperating may or may not be their first choice. Thus, the observed outcome is the result of the joint multilateral interaction between a player and that player's potential partners, as neither the player nor any one partner can unilaterally determine the cooperation outcome.

In summary, the new model is suited to analyze scenarios where (i) a player is confronted with multiple, mutually exclusive choice alternatives, and (ii) the player's potential partners can opt in or out of a cooperation with that player. Because the alternatives are mutually exclusive, the model is designed to examine substitution not complementary effects. To better outline the scope conditions of the new estimator, let us consider three scenarios that involve an interaction among three (or more) actors: direct and mediated multilateral interactions, and bilateral interactions with third party intervention.²

A direct multilateral event is, for example, where three countries come together and negotiate a treaty, of which all three can be a part. When the available choices are substitutes, though, only one of the competing alternatives can be chosen. This is the type of interaction the new estimator is designed to model. For example, one negotiates with two potential employers, and, if both make an offer, the employee must choose one. This is still a multi-actor interaction; the potential employers interact, albeit indirectly via the future employee. Hearing that employer *B* made a bigger offer, employer *A* can decide to match, or raise the offer. In fact, both employers can make offers at the same time, and may do so either in response or

¹ In this paper I use the terms (empirical) model, estimator, and estimation procedure interchangeably.

² I thank an anonymous reviewer for pointing out these scenarios.

anticipation of the other. I refer to this scenario as a multilateral mediated interaction.

Lastly, in bilateral events subject to third party influence, the intervening third party does not represent a plausible alternative to either of the remaining two actors. This would be the case of a third country promising military support to a government engaged in a civil conflict. The government, which otherwise would have signed a peace agreement with the rebels, now decides to continue fighting. As the external support is meant to augment the regime's military forces, the third party does not constitute a distinct fighting opponent for the rebel group, and hence not a proper alternative to the government.

In the next section, I introduce the new MMIwPO model. To validate the new estimator, I then present MC simulation results which provide strong evidence for the superior performance of the MMIwPO relative to the standard approach. I also analyze real data on defense alliance formation, and then compare the MMIwPO results to those reported by a univariate logit to see what would be different were we to use the standard methodology in the literature (Gibler & Wolford 2006, Poast 2010). Finally, for researchers interested in employing the MMIwPO, I provide detailed working examples, as well as an easy-to-use program that implements the estimation procedure in Stata.

2 Analyzing Multiparty Negotiations

The current practice of analyzing multi-actor interaction as the (joint) binary choice of either a single or two decision-makers is not optimal for both theoretical and methodological reasons. Theoretically, scholars are often interested in the motivations of all parties involved, as well as in how the decision of one actor constrains or alters the actions of others. Yet, univariate models can only provide insights into how factors affect the incentives (probability) of one single actor (process), while bivariate models can accommodate interactions between only two actors. Methodologically, analyzing multilateral processes as a set of either monadic or dyadic events only increases the sample size, but it adds no new information and introduces bias (Croco & Teo 2005, Fordham & Poast 2016, Poast 2010, Signorino 1999). To account for these challenges, an ideal estimator should be able to accommodate the input of multiple actors into a unified, overarching decision making process.

2.1 The Multilateral Mediated Interaction with Partial Observability Model

Multilateral negotiations can be represented as an interaction between two (or more) sets of actors, $\{1, 2, \dots, N\}$ and $\{1, 2, \dots, K\}$, where at least of one set comprises minimum two elements.³ Each actor $i \in N$ must decide whether or not they want to cooperate with every actor $j \in K_i$, and vice versa. For each actor i , we can write the net value of the cooperation with potential partner j as an unobserved, latent variable

$$y_{ij}^* = x_{ij}\beta + \epsilon_{ij} \quad (1)$$

where β represents a vector of coefficients, x_{ij} represents a vector of independent variables for actor i associated with potential partner j , and ϵ_{ij} is assumed to be drawn from the standard logistic distribution with a mean of 0 and a variance of $\frac{\pi^2}{3}$.

Actor i is willing to cooperate with j if the net value of cooperation is positive:

$$y_{ij} = \begin{cases} 1 & \text{if } y_{ij}^* > 0 \\ 0 & \text{if } y_{ij}^* \leq 0. \end{cases} \quad (2)$$

Similarly, for each actor j , we can describe the net value of the cooperation with potential partner i

³ Were both sets a singleton, the scenario reduces to a bilateral interaction. Theoretically, though, the process underpinning both bilateral and multilateral interactions is the same. Fordham and Post's (2016, 841) assessment of the alliance formation process can be easily generalized to encompass multi-actor interactions generally: "Scholars should reconceptualize *all* alliances, bilateral and multilateral, as originating from a multilateral process. Even when states only form a bilateral alliance, the alliance, in principle, could have contained additional states". In the context of bilateral negotiations, some outside negotiations can be observed only in hindsight, as it was the case in the E.U.-Ukraine interaction. Also, it may be the case that negotiations with a potential partner never even start when an actor infers that the negotiations would fail. Yet, the actor had actually weighed the opportunity to engage in that negotiation. The decision to turn down that opportunity is equivalent to rejecting the potential partner.

as an unobserved, latent variable

$$y_{ji}^* = z_{ji}\gamma + \epsilon_{ji} \quad (3)$$

where γ represents a vector of coefficients, z_{ji} represents a vector of independent variables for actor j associated with potential partner i , and ϵ_{ji} is assumed to be drawn from the standard logistic distribution.

As with the actor i , actor j is willing to cooperate with i if the net value of cooperation is positive:

$$y_{ji} = \begin{cases} 1 & \text{if } y_{ji}^* > 0 \\ 0 & \text{if } y_{ji}^* \leq 0. \end{cases} \quad (4)$$

If the unobserved variables influencing actors i and j are independent, $\text{cov}(\epsilon_{ij}, \epsilon_{ji}) = 0$,⁴ we can write the individual probability that i wants to cooperate with j as

$$\Pr(y_{ij} = 1 | x_{ij}) = P_{ij} = F(x_{ij}\beta), \quad (5)$$

and the individual probability that j wants to cooperate with i as

$$\Pr(y_{ji} = 1 | z_{ji}) = P_{ji} = F(z_{ji}\gamma) \quad (6)$$

where $F(\cdot)$ is the cumulative distribution function of the standard logistic distribution, i.e., $F(\cdot) = \Lambda(\cdot)$.

Yet, when it comes to bilateral interactions we often do not observe the individual choices of the actor i , y_{ij} , or j , y_{ji} , only the *joint* outcome $o_{ij} = y_{ij} \times y_{ji}$. o_{ij} equals one only if y_{ij} and y_{ji} are both one, zero otherwise. Where this is the case, the choices of actors are only *partially* observed (Abowd & Farber 1982, Poirier 1980, Przeworski & Vreeland 2002, Nieman 2015). On the one hand, it is possible to infer the choices of both actors when they cooperate – actor i must have chosen to cooperate with j , and actor j must have agreed. On the other hand, it is not possible to infer the choices of either actor when they do not cooperate. The failure to observe i and j cooperating could be because actor i did not want to cooperate with j , actor j rejected i , or neither actor was interested in a mutual cooperation. Therefore, the

⁴ I relax this assumption and discuss the ensuing consequences in Subsection 2.2.

joint probability that actor i seeks to cooperate with j *and* that actor j accepts to cooperate with i is

$$\Pr(o_{ij} = 1 | x_{ij}, z_{ji}) = P_{ij} \times P_{ji}. \quad (7)$$

Conversely, the probability that actor i does not cooperate with actor j , either because i refuses to cooperate with j or because j rejects i (or both), is simply

$$\Pr(o_{ij} = 0 | x_{ij}, z_{ji}) = 1 - (P_{ij} \times P_{ji}). \quad (8)$$

Moving from bilateral to multilateral negotiations, the situation is further complicated by the fact that, for actor i , the (i, j) dyad is not the only cooperation opportunity. In fact, actor i is confronted with a multichotomous choice, not a dichotomous one. Thus, the above joint (partial observability) probabilities, Eq. (7) and Eq. (8), apply to every single cooperation opportunity. Clearly, for actor i , there are K_i such opportunities. How can we then compute the probability that actor i cooperates with actor j , and not with $\sim j$, out of a set of K_i potential partners?⁵ Starting from the joint probabilities, we need to compute the following *conditional* probability

$$\begin{aligned} & \Pr(y_{ij} = 1, y_{i2} = 0, \dots, y_{iK_i} = 0 \mid y_{ij} + y_{i2} + \dots + y_{iK_i} = 1) = \\ & \frac{\Pr(y = (1, 0, \dots, 0)')}{\Pr(y = (1, 0, \dots, 0)') + \Pr(y = (0, 1, \dots, 0)') + \dots + \Pr(y = (0, 0, \dots, 1)')} \end{aligned} \quad (9)$$

where $j \in K_i$.

Using the joint dyadic probabilities in Eq. (7) and Eq. (8), the numerator in Eq. (9) can be computed

⁵ In line with the real life examples that motivate this paper, one assumption of the model developed here is that actor i can choose *and* be accepted by only one actor from the choice set K_i . The reason for this is that the potential partners are presumed to be substitutes not complements.

as

$$\begin{aligned} & \left\{ [P_{ij}P_{ji}] \times [1 - P_{i2}P_{2i}] \times \dots \times [1 - P_{iK_i}P_{K_i i}] \right\} = \\ & \left\{ \left[\left(\frac{e^{x_{ij}\beta}}{1 + e^{x_{ij}\beta}} \right) \left(\frac{e^{z_{ji}\gamma}}{1 + e^{z_{ji}\gamma}} \right) \right] \times \left[1 - \left(\frac{e^{x_{i2}\beta}}{1 + e^{x_{i2}\beta}} \right) \left(\frac{e^{z_{2i}\gamma}}{1 + e^{z_{2i}\gamma}} \right) \right] \times \dots \times \left[1 - \left(\frac{e^{x_{iK_i}\beta}}{1 + e^{x_{iK_i}\beta}} \right) \left(\frac{e^{z_{K_i i}\gamma}}{1 + e^{z_{K_i i}\gamma}} \right) \right] \right\}, \end{aligned} \quad (10)$$

and the denominator as

$$\begin{aligned} & \left\{ [P_{ij}P_{ji}] \times [1 - P_{i2}P_{2i}] \times \dots \times [1 - P_{iK_i}P_{K_i i}] \right\} + \left\{ [1 - P_{ij}P_{ji}] \times [P_{i2}P_{2i}] \times \dots \times [1 - P_{iK_i}P_{K_i i}] \right\} + \\ & \dots + \left\{ [1 - P_{ij}P_{ji}] \times [1 - P_{i2}P_{2i}] \times \dots \times [P_{iK_i}P_{K_i i}] \right\} = \\ & \left\{ \left[\left(\frac{e^{x_{ij}\beta}}{1 + e^{x_{ij}\beta}} \right) \left(\frac{e^{z_{ji}\gamma}}{1 + e^{z_{ji}\gamma}} \right) \right] \times \left[1 - \left(\frac{e^{x_{i2}\beta}}{1 + e^{x_{i2}\beta}} \right) \left(\frac{e^{z_{2i}\gamma}}{1 + e^{z_{2i}\gamma}} \right) \right] \times \dots \times \left[1 - \left(\frac{e^{x_{iK_i}\beta}}{1 + e^{x_{iK_i}\beta}} \right) \left(\frac{e^{z_{K_i i}\gamma}}{1 + e^{z_{K_i i}\gamma}} \right) \right] \right\} + \\ & \left\{ \left[1 - \left(\frac{e^{x_{ij}\beta}}{1 + e^{x_{ij}\beta}} \right) \left(\frac{e^{z_{ji}\gamma}}{1 + e^{z_{ji}\gamma}} \right) \right] \times \left[\left(\frac{e^{x_{i2}\beta}}{1 + e^{x_{i2}\beta}} \right) \left(\frac{e^{z_{2i}\gamma}}{1 + e^{z_{2i}\gamma}} \right) \right] \times \dots \times \left[1 - \left(\frac{e^{x_{iK_i}\beta}}{1 + e^{x_{iK_i}\beta}} \right) \left(\frac{e^{z_{K_i i}\gamma}}{1 + e^{z_{K_i i}\gamma}} \right) \right] \right\} + \\ & \dots + \\ & \left\{ \left[1 - \left(\frac{e^{x_{ij}\beta}}{1 + e^{x_{ij}\beta}} \right) \left(\frac{e^{z_{ji}\gamma}}{1 + e^{z_{ji}\gamma}} \right) \right] \times \left[1 - \left(\frac{e^{x_{i2}\beta}}{1 + e^{x_{i2}\beta}} \right) \left(\frac{e^{z_{2i}\gamma}}{1 + e^{z_{2i}\gamma}} \right) \right] \times \dots \times \left[\left(\frac{e^{x_{iK_i}\beta}}{1 + e^{x_{iK_i}\beta}} \right) \left(\frac{e^{z_{K_i i}\gamma}}{1 + e^{z_{K_i i}\gamma}} \right) \right] \right\}. \end{aligned} \quad (11)$$

Upon simplifying both expressions,⁶ the conditional probability that actor i cooperates with actor j , out of a set of K_i alternatives, can be written as

$$\Pr(y_{ij} = 1 | x_{i1}, \dots, x_{iK_i}, z_{1i}, \dots, z_{K_i i}; \sum_{k=1}^{K_i} y_{ik} = 1) = \frac{\frac{e^{x_{ij}\beta} e^{z_{ji}\gamma}}{1 + e^{x_{ij}\beta} + e^{z_{ji}\gamma}}}{\sum_{k=1}^{K_i} \frac{e^{x_{ik}\beta} e^{z_{ki}\gamma}}{1 + e^{x_{ik}\beta} + e^{z_{ki}\gamma}}} \quad (12)$$

where $j \in K_i$.

⁶ In Online Appendix A I present step-by-step the simplification process.

It follows that the likelihood for the entire sample in the MMIwPO model is

$$\mathcal{L} = \prod_{i=1}^N \prod_{k=1}^{K_i} P_{ik}^{d_{ik}}, \quad (13)$$

where d_{ik} is an indicator variable that equals one if actor i cooperates with actor k , zero otherwise.

As with all multivariate models, in order for the MMIwPO to be identified, it is necessary that there be at least one variable in one of the vectors of independent variables (x_{ij} , z_{ji}), but not the other (Poirier 1980, Gordon & Smith 2004).

A Multiparty Interaction Example

An example might help to illustrate how actors' partially observed preferences determine the observed outcome, as well as the data structure underlying a multi-actor interaction scenario. Suppose we have a set of N states and a set of K military alliances. The question at hand is which alliance a given candidate state (e.g., state $S_1 \in N$) joins, given that it can be a member of only one alliance. In this particular example, S_1 is confronted with five alliance alternatives: A_1 , A_2 , A_3 , A_4 , and A_5 , where $A_* \in K$. Since theoretically a candidate state may affiliate with any alliance, the data are organized into alliance membership opportunity sets comprising all state-alliance dyads. Specifically, the alliance membership opportunity set of state S_1 comprises five state-alliance dyads: $S_1 - A_1$, $S_1 - A_2$, $S_1 - A_3$, $S_1 - A_4$, and $S_1 - A_5$.

To each state-alliance dyad corresponds a pair of independent variable sets: one for the candidate state, $x_{S_*A_*}$, and the other one for the target alliance, $z_{A_*S_*}$. Whether or not a given actor cooperates is a function of that actor's determinants. Facing multiple alternatives, the state ranks alliances in terms of the utility it receives from membership in that alliance. In this example, S_1 ranks the five alliances as follows. Ideally, it would like to be a member of A_1 , but, if this is not possible, it is content to join A_3 . If everything else fails, the candidate state would also settle for A_2 . However, it would join neither A_4 nor A_5 under any circumstances. For their part, each alliance decides whether to accept or reject the wouldbe member. Apart from A_1 and A_5 , all alliances would welcome S_1 . The preferences of all actors and the resultant outcome are shown in Table 1. Note that the information presented in last three columns, to the right of the vertical line, are not available to the researcher.

Table 1: Alliance Membership as Multilateral Negotiations

Candidate State	Target Alliance	Candidate Determinants	Alliance Determinants	Observed Joint Outcome $y_{S^*A^*}$	Candidate Rank Choice	Unobserved Individual Preferences	
						y_{S^*}	y_{A^*}
S_1	A_1	$x_{S_1A_1}$	$z_{A_1S_1}$	0	1	1	0
S_1	A_2	$x_{S_1A_2}$	$z_{A_2S_1}$	0	3	1	1
S_1	A_3	$x_{S_1A_3}$	$z_{A_3S_1}$	1	2	1	1
S_1	A_4	$x_{S_1A_4}$	$z_{A_4S_1}$	0	—	0	1
S_1	A_5	$x_{S_1A_5}$	$z_{A_5S_1}$	0	—	0	0
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
S_N	A_{k_N}
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
S_N	A_{K_N}

Note: S^* is an indicator for candidate states, where $S^* \in N$. A^* is an indicator for target alliances, where $A^* \in K$. $y_{S^*A^*}$ represents the observed outcome. $x_{S^*A^*}$ and $z_{A^*S^*}$ represent the candidate state and target alliance determinants, respectively. y_{S^*} represents the unobserved preferences of state S^* with respect to whether it is willing to seek membership with a given alliance. y_{A^*} represents the unobserved preferences of target alliances with respect to whether they are interested to accept that state.

Which alliance does state S_1 ultimately join? In this example, S_1 becomes a member of A_3 because its most preferred choice, A_1 , rejects the state, and among the remaining viable choices, A_2 and A_3 , it prefers A_3 the best. This hypothetical six-party interaction highlights the need to employ a multivariate partial observability model in a multilateral negotiation scenario. In the current set-up there are several substantively different, yet observationally equivalent, non-cooperation outcomes. Specifically, S_1 joins none of the following alliances: A_1 (the state’s first choice), A_2 (an alliance that S_1 actually finds acceptable, and vice versa), A_4 (an alliance the state rejects), and A_5 (an alliance the state rejects, and vice versa). As Poast (2012, 279) argues, there is the value in considering the “dogs that didn’t bark” observations, and partial observability models are designed to make the most of these non-cooperation outcomes.

For example, like bivariate partial observability models, MMIwPO can accommodate situations where non-cooperation is the result of one actor rejecting the other (e.g., the $S_1 - A_1$ and $S_1 - A_4$ interactions), or both parties rejecting cooperation (e.g., the $S_1 - A_5$ interaction). Unlike the standard partial observability models, however, MMIwPO can also explain why a cooperation that is mutually beneficial to both parties, such as the interaction between S_1 and A_2 , does not occur. Recall that in the MMIwPO model, the utility associated with a specific alliance is evaluated *relative* to that of the other alternatives. The fact

that another alternative is comparatively more advantageous can explain why a bilateral cooperation, whose *absolute* value is positive, does not occur.

As a side note, in this example, the set of five potential partners was introduced without a preliminary discussion. Yet, determining the pool of alternatives to include in the estimation is a very important issue. When it comes to identifying the full set of possible alternatives, the research question should inform the selection criteria. For example, if the analysis is about interstate conflict, then all other states are potential targets or aggressors. Often, though, only a small subset of all dyads can plausibly experience the phenomenon in question (e.g., the odds of a war between Saint Kitts and Nevis and San Marino are slim to none). Some researchers argue that only these dyads are a proper reference group, and the others should be discarded (Lemke & Reed 2001).

When this is a concern, one could run a global goodness of fit test to determine whether adding a given alternative improves the overall model fit. This approach is advised only if the alternative in question is included in all (or most) actors' pool of potential partners. To identify the subset of relevant dyads at the individual dyad level, one could run a separate logit model and then include in the MMIwPO only the dyads with a high enough likelihood. This approach, however, does not properly account in the final estimation for the uncertainty in identifying these dyads. Alternatively, one can employ theoretical criteria to tag relevant dyads. A popular approach in the conflict literature, for example, is to restrict the analysis to "politically relevant" dyads, usually defined in terms of the dyad members' major power status and geographic contiguity (Lemke & Reed 2001). While the precise attributes are likely to vary with the research question, theory should play a central role in informing this decision.

2.2 The Multilateral Mediated Interaction with Partial Observability and Odds Ratio

The joint probability, Eq. (7), and the subsequent conditional joint probability discussed above, Eq. (12), assume that the two component probabilities, P_{ij} and P_{ji} , are independent. Of course, in some scenarios this assumption might not be warranted. Where this is the case, we must address the problem that both the systematic and/or stochastic components could be correlated (Leemann 2014, Poirier 1980). In practice, we need to define another parameter, the odds ratio, that captures the relationship between the two outcomes. The odds ratio, ψ , is defined as $\psi = \frac{\Pr(P_{ij}=1 \mid P_{ji}=1) \times \Pr(P_{ij}=0 \mid P_{ji}=0)}{\Pr(P_{ij}=1 \mid P_{ji}=0) \times \Pr(P_{ij}=0 \mid P_{ji}=1)}$, where $\psi = 1$ if and only if the

individual probabilities are independent (Bland & Altman 2000, Szumilas 2010). Solving the equation allows us to compute the joint outcome, $P_{o_{ij}}$, as a function of the odds ratio, ψ , and the marginal individual probabilities, P_{ij} and P_{ji} . The solution is a quadratic equation in $P_{o_{ij}}$ ⁷

$$\begin{aligned} 0 &= a \times P_{o_{ij}}^2 + b \times P_{o_{ij}} + c \\ P_{o_{ij}} &= \left(-b + \sqrt{(b^2 - 4 \times a \times c)} \right) / (2 \times a) \end{aligned} \quad (14)$$

where $a = 1 - \psi$, $b = 1 - (1 - \psi) \times (P_{ij} + P_{ji})$, and $c = -\psi \times P_{ij} \times P_{ji}$ (McLerran 2010).

When $\psi = 1$, then $a = 0$, and therefore the denominator of Eq. (14) is zero. Consequently, we cannot compute the $P_{o_{ij}}$ employing the above formula. However, recall that when $\psi = 1$, it must be the case that the individual probabilities are independent, so the joint probability in this scenario is simply the product of the two probabilities. Therefore, the new joint probability that actor i seeks to affiliate with actor j , and that j accepts i is

$$\Pr(o_{ij}|x_{ij}, z_{ji}; \psi = 1) = \begin{cases} \left(-b + \sqrt{(b^2 - 4 \times a \times c)} \right) / (2 \times a) & \text{if } \psi \neq 1; \\ P_{ij} \times P_{ji} & \text{if } \psi = 1. \end{cases} \quad (15)$$

Substituting the new joint outcome probability in the conditional probability defined in Eq. (9), we obtain the following conditional probability that actor i cooperates with actor j , out of a set of K_i potential partners,

$$\Pr(y_{ij} = 1 | x_{i1}, \dots, x_{iK_i}, z_{1i}, \dots, z_{K_i i}; \sum_{k=1}^{K_i} y_{ik} = 1) = \frac{\prod_{k=1}^{K_i} \left[P_{o_{ik}}^{y_{ik}} (1 - P_{o_{ik}})^{1-y_{ik}} \right]}{\sum_{s_i \in S_i} \prod_{k=1}^{K_i} \left[P_{o_{ik}}^{s_{ik}} (1 - P_{o_{ik}})^{1-s_{ik}} \right]}, \quad (16)$$

where y_{ik} is the dependent variable taking on values 0 or 1, $y_i = (y_{i1}, \dots, y_{iK_i})$ is the outcomes for actor i as a whole, s_{ik} is equal to 0 or 1 with $\sum_{k=1}^{K_i} s_{ik} = 1$, and S_i is the set of all possible combinations of y_i one and zeros. Clearly, there are K_i such combinations.

⁷ In Online Appendix B I present the intermediate steps in solving the equation.

The Multilateral Mediated Interaction with Partial Observability vs. Conditional Logit

How does the new MMIwPO model relate to extant choice models? It turns out that MMIwPO nests the conditional logit (CL), the standard univariate model for multiple choices (McFadden 1974, McFadden 1974a).⁸ In the restricted model, the CL, a player chooses among multiple choice alternatives and the player's decision is decisive. In effect, the outcome is determined solely by the decision of a single actor. In the MMIwPO, the premise is similar: a player chooses among multiple potential partners. The twist is that the partners have a say in the final outcome, as they can reject individual requests for cooperation. Thus, the outcome in this scenario is determined by the joint choices of a player and that player's potential partners.

The fact that the CL is only a special case of the unrestricted MMIwPO model means that, on the one hand, the two models produce identical results in situations where the outcome is indeed the result of a single actor's unilateral decision. On the other hand, the MMIwPO ought to produce more accurate estimates in situations where the outcome is the result of multilateral interactions. To determine whether this is indeed the case, I conduct multiple MC experiments that allow me to directly compare the performance of the MMIwPO and CL models. MC simulations are useful because they provide a controlled environment, where one can tease out the effect of a given modeling choice. To prevent contamination effects, ideally, there should be just one difference between the compared scenarios. The only difference between the two models in question is that the new estimator allows potential partners to reject individual requests for cooperation. In fact, were we change the corresponding probabilities to reflect that all K actors automatically accept player

⁸ The MMIwPO model is also similar to the two-sided logit (TSL) that has been developed to examine how the choices of workers and employers jointly determine employment outcomes (Logan 1996, Logan 1998). There are at least three notable differences between the TSL and the MMIwPO model. First, the likelihood function associated with the TSL model is not globally concave. This causes practical problems because its estimation algorithm often finds a local not a global maximum (Logan 1996, 127). Second, the TSL restricts a player's choice to the subset of partners that are willing to cooperate. This assumes that we can always observe or infer the individual choices of all actors. Third, the TSL cannot estimate whether actors' choices are interrelated. This is due in part to the fact that the model presupposes that actors move sequentially, and therefore the disturbances are assumed uncorrelated. The MMIwPO does not suffer from any of these problems.

i (i.e., $\frac{\exp(z_{ki}\gamma)}{1+\exp(z_{ki}\gamma)} = 1$ in Eq. (10) and (11)), the MMIwPO probability simplifies to the CL's conditional probability. Thus, when comparing the MMIwPO and CL estimates, it is easy to tell what drives the changes in results, if any.

3 Monte Carlo Simulations

Each MC experiment comprises 2,000 players and 5 potential partners. Each player chooses whether they want to cooperate with any, or none, of the five alternatives. Similarly, each potential partner decides whether they would reject or accept a player were he willing to cooperate. Both players and potential partners choose the option that maximizes the utility accrued from that cooperation. Since there are two sets of actors, i.e., the players (N) and potential partners (K), there are two distinct utility functions:

$$y_{ij}^* = \beta_1 IV1_{ij} + \beta_2 IV2_{ij} + \epsilon_{ij} \quad (17)$$

$$y_{ji}^* = \gamma_1 IV1_{ji} + \gamma_2 IV3_{ji} + \epsilon_{ji} \quad (18)$$

where $i \in N, j \in K, \beta_1 = 0.2, \beta_2 = -0.6, \gamma_1 = +/ - 0.8, \gamma_2 = 0.4$, and $\text{cov}(\epsilon_{ij}, \epsilon_{ji}) = \rho$.

There are two inputs to both players and potential partners' utilities. One of the inputs, the independent variable $IV1$, is the same in both actors' utility function. This is to evaluate how a factor can simultaneously affect the incentives of both players and partners. The $IV1$'s coefficient β_1 is always positive. In contrast, while retaining the same absolute value, γ_1 is either positive or negative in alternative MC experiments. This allows us to evaluate how different estimators perform when the common factor has either a similar (positive in this case) or opposing effect on the actors. The second input is distinct between the two utilities; $IV2$ is a determinant of the player's incentives alone, whereas $IV3$ is included only in the potential partners' equation. All three independent variables have values randomly drawn from the standard logistic distribution, with a mean of 0 and a variance of $\frac{\pi^2}{3}$. The error terms are randomly drawn from a cumulative logistic distribution with a correlation coefficient of ρ .

Next, employing the logistic function, I compute two dichotomous variables to indicate the probabil-

ity that actor i is willing to cooperate with j , and vice versa:

$$y_{ij} = \begin{cases} 1 & \text{if } \frac{1}{1+e^{-y_{ij}^*}} > 0.5 \\ 0 & \text{if } \frac{1}{1+e^{-y_{ij}^*}} \leq 0.5; \end{cases} \quad (19)$$

$$y_{ji} = \begin{cases} 1 & \text{if } \frac{1}{1+e^{-y_{ji}^*}} > 0.5 \\ 0 & \text{if } \frac{1}{1+e^{-y_{ji}^*}} \leq 0.5. \end{cases} \quad (20)$$

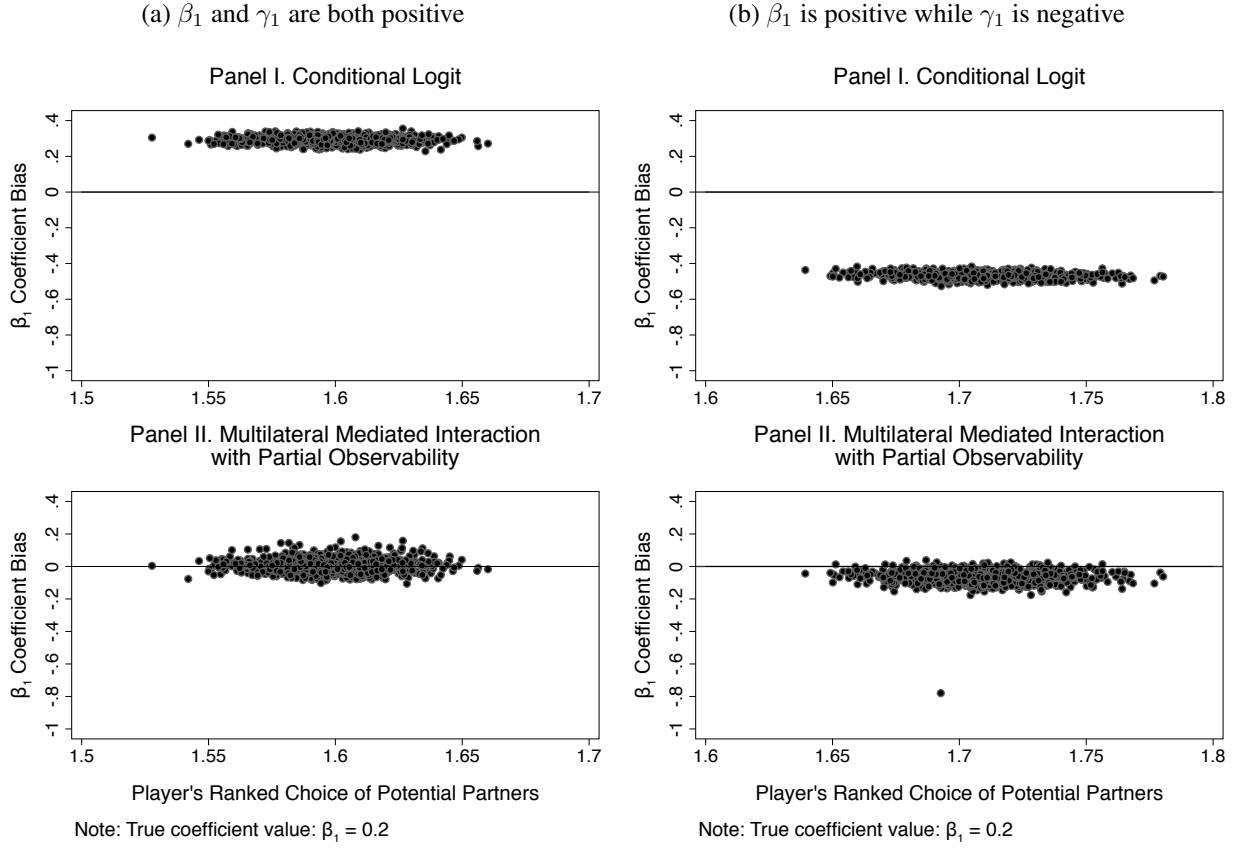
One assumption of the MMIwPO model is that potential partners are substitutes not complements. In practical terms this means that actor i can choose and be accepted by only one actor from the choice set K_i . Therefore, the dependent variable is coded one for the (y_{ij}^*, y_{ji}^*) dyad that maximizes the product of the two latent variables, given that both individual probabilities are one. Formally, the dependent variable, y_{ij} , equals one if $\left((y_{ij}^* \times y_{ji}^*) = \max_{j \in K_i} \left\{ (y_{i1}^* \times y_{1i}^*), \dots, (y_{iK_i}^* \times y_{K_i i}^*) \right\} \mid (y_{ij} = 1 \mid y_{ji} = 1) \right)$, zero otherwise.

Overall, I conduct thirty MC experiments, which represent the combination between (i) five different values for the disturbance correlation parameter (ρ successively takes the value of 0, 0.25, 0.5, 0.75, and 1), (ii) two different values for the γ_1 coefficient ($+/-0.8$), and (iii) three different estimators (CL, MMIwPO, and MMIwPO with Odds Ratio). For each of the thirty scenarios, I run 1,000 simulations. To keep things concise, in the main text I present only the results from the CL and MMIwPO models where the disturbances are not correlated, $\rho = 0$.⁹ Since the CL can accommodate only one set of actors, the MC experiments in this case are restricted to analyzing the players' incentives to choose among available potential partners. Thus, right off the bat, the univariate logit fails to provide any insights into how factors affect the potential partners' incentives to accept or reject the players' (potential) cooperation requests.

Which estimator performs best according to the MC simulations? I employ two indicators to discriminate between the MMIwPO and CL models. First, I perform a likelihood-ratio test to assess the goodness-of-fit between the two empirical models. Without exception, the test indicates that the MMIwPO model fits the simulated data significantly better in every MC experiment. Thus, the MMIwPO's increased complexity is justified in terms of the significant improvement in fit over the CL model. Second, I contrast the MMIwPO and CL models in terms of their ability to minimize the bias in the estimated coefficients.

⁹ In Online Appendix C I present and discuss the full set of MC experiments.

Figure 1: The β_1 Coefficient Bias



Note: Figure 1 illustrates the β_1 coefficient bias for the CL and MMIwPO models. The bias estimates are based on 1,000 MC simulations. The left panels indicate scenarios where the common variable, $IV1$, has a positive effect on both sets of actors, N and K . The right panels indicate scenarios where $IV1$ has a positive effect on the N actors, and a negative effect on the K s.

Figure 1 shows the bias in the β_1 coefficient across the player's ranked choice of potential partners for both the CL and MMIwPO models.¹⁰ The coefficient bias is simply the difference between the estimated and the true value of that coefficient. The ranked choice captures potential partners' position on the player's preference scale. For example, a value of one indicates that the player was accepted by the partner that he ranks first. Higher values indicate that the player was rejected by his preferred choice, and he had to settle for his second best alternative, or worse. The left panels indicate scenarios where the common factor, $IV1$, has a positive effect on both sets of actors, N and K . The right panels indicate scenarios where $IV1$ has a

¹⁰ In the main text I discuss only the bias in the estimated coefficient on the variable present in both actors' utility function, $IV1$. However, I do report the estimated bias for all four coefficients in Online Appedix C.

positive effect on the N actors, and a negative effect on the K s.

It is easy to note that in both scenarios the MMIwPO reports the least biased coefficients, with the bias values straddling the zero line (see Panel II in Figure 1a and 1b). While the CL's β_1 estimates are always biased, the direction of the bias is determined by the sign of the γ_1 coefficient from the potential partners' utility function. Recall that β_1 is always positive. When γ_1 is also positive, the estimated β_1 coefficient is always bigger than its true value since it is artificially pulled up by the positive effect that $IV1$ has on the potential partners' incentives (Figure 1a Panel I). When γ_1 is negative, the CL's β_1 coefficient is always smaller since it is pulled down by $IV1$'s negative effect on the partners' incentives (Figure 1b Panel I). In effect, the conditional logit model reports a weighted average of the $IV1$'s effect on all actors.

4 An Application to Alliance Formation

In this section, I demonstrate the inferential benefits of employing the MMIwPO by applying the model to real data on alliance formation. The upcoming empirical analysis is meant only to illustrate the applicability of the MMIwPO model, not to develop a new theoretical account of the alliance formation process. First, I present a couple of anecdotes to support the claim that a state's membership in a given alliance is contingent on the available options. Next, I describe the data and address various measurement issues. I then present the results and discuss their substantive importance.

In June 2013, the Colombian President Juan Manuel Santos announced that he wants a closer partnership with NATO. Not surprisingly, his announcement stirred strong emotions. Colombia's partners in the Union of South American Nations, the defense alliance in South America, expressed strong objections to this proposal with the Bolivian President claiming that "[a]ny presence of NATO in South America or Latin America poses a threat to peace in the region" (Mallén 2013). In order to contain the problem and appease Colombia's current South American allies, the Colombian foreign minister qualified the President's statement by saying that Colombia is not actively seeking to join NATO (Mallén 2013). NATO also distanced itself from the Colombian President's remarks saying that Colombia cannot become a NATO member as it "does not meet the geographically limited membership criteria" (Agence France-Presse 2013). Of course, there are various reasons why an alliance might find a particular candidate state undesirable. Georgia, for example, initiated the NATO membership process in 1994. After more than twenty years, it is still not a

member because of political and military considerations (Fuller 2014). Indeed, it may be the case that negotiations over a potential alliance membership never even start because the interested candidate knows that it does not meet the necessary requirements, or it infers that the negotiations would fail.

Regardless of whether because of a geographic or other technical reason, when states are denied membership in an alliance of their choice, they are forced to look for security guarantees somewhere else. Therefore, their subsequent decision (i.e., to join another alliance, form a new one, or remain neutral) is determined in part by having had their options curtailed by their preferred alliance's rejection. For example, Fordham and Poast (2016) argue that in 1963 Spain was forced to form a bilateral alliance with the U.S., instead of joining the existing U.S. sponsored European alliance, because of other NATO members' distaste for Franco. In their assessment, this "illustrates how the positions of third-party states can influence the formation of a bilateral alliance" (Fordham & Poast 2016, 844).

4.1 Data and Measurement Issues

The upcoming analysis builds on Poast's (2010) work on alliance formation, which in turn draws on Gibler and Wolford's (2006) study on alliance membership. While my analysis closely follows Poast (2010) with respect to the set of explanatory variables, operationalization rules, and temporal domain, it is not a replication *per se* of that work. The two studies differ in terms of their scope and the cases included in the respective analyses. On the one hand, Poast (2010) addresses the question of why a given k -adic alliance forms while alternative k -member alliances do not. On the other hand, this study answers two questions: (i) why a state forms a given k -adic alliance instead of joining an existing one, and (ii) why the other alliance members accept that state. Finally, Poast (2010) drops alliances with six or more members from the analysis, whereas I include all alliances irrespective of their size. These differences aside, I focus on Poast's (2010) study because, by conceptualizing alliance formation as a multilateral event, it departs from the traditional dyadic analysis framework. In the process it also provides the first guidelines with respect to various aggregation issues that one faces when moving from a dyadic to multi-actor interaction.

Following Poast (2010), the analysis focuses solely on alliance formation, treating the decision to join an existing alliance as distinct from the decision to create a new one.¹¹ The data are organized into

¹¹ In this study I look only at defense alliances. One assumption of the MMIwPO model is that the

alliance membership opportunity sets. A state's membership opportunity set comprises the state-alliance dyads between the potential member and all extant alliances, as well as the dyad between that state and the new alliance.¹² An example might help to illustrate the data structure. Suppose we have a state, S , that forms a new alliance, A_3 , at a time when there were two other alliances, A_1 and A_2 . In this scenario, the alliance membership opportunity set of state S would comprise three state-alliance dyads: $S - A_1$, $S - A_2$, and $S - A_3$. Including existing alliances in a state's membership opportunity set captures the intuition that, aside from forming a new alliance, the state also has the option to reach out to extant alliances. Indeed, the need to form a new military alliance can stem from the fact that either no existing alliance is able to fulfill that state's security needs, or, if such alliances exist, they reject that state (e.g., Spain in 1963, Colombia and Georgia nowadays). In this sense, the new alliance is a substitute to existing alliances.¹³

The temporal domain covers the period from 1816 till the end of the cold war in 1990. The number of independent states and military alliances varies across time, but the full dataset comprises 121 states and 182 military alliances. In total, there are 5,638 unique state-alliance dyads grouped in 509 alliance membership opportunity sets.¹⁴ The dependent variable in the upcoming empirical analysis, *New Allies*, is coded one if available alternatives are potential substitutes. Given this assumption, it does not make sense to put defense and entente pacts together. Previous studies have also recognized this important distinction and analyzed defense alliances, which entail the highest level of military commitment, as a stand-alone category (Gibler & Wolford 2006, Lai & Reiter 2000).

¹² In a state's membership opportunity set, any new alliance is coded as being formed from all other members except the candidate state. This rule facilitates the comparison between the new alliance and those alliances of which the state never becomes a member. For the same reason, when computing various alliance specific variables, the candidate state's input is ignored for all alliances in that state's membership opportunity set.

¹³ It is worth noting that alternative research questions may emphasize the potential complementarity among alliances. This is generally the case of analyses that examine states' membership in multiple organizations that are similar in scope. Analyses that focus on complementary effects, though, are not fully isomorphic with the assumptions underlying the new estimator.

¹⁴ For the cases where a state forms more than one alliance in a given year, I create duplicates of the

the state-alliance dyad represents a true relationship in the sense that the state is an original member of the newly formed alliance, zero otherwise.

In Poast's (2010) analysis, four core determinants explain a state's decision to form a new alliance: capability ratio, joint democracy status, geographic distance, and common threat. Since the main objective of a defense alliance is to provide assistance in case of foreign aggression, the military capabilities of the parties involved play a crucial role in determining whether a state joins a given alliance. *Capability Ratio* captures the relative military strength of the candidate state and target alliance. Effectively, it is an indicator of whether the wouldbe member is likely to be a consumer or provider of security within a particular alliance. *Capability Ratio* is calculated as the ratio between the military capability of the wouldbe member and that of the target alliance.¹⁵ For example, if C_S and C_A represent the military capabilities of the candidate state and target alliance, respectively, the *Capability Ratio* would be calculated as: $\frac{C_S}{C_S + C_A}$. On the one hand, individual states are more likely to seek membership with a strong military alliance that can protect them from outside threats. On the other hand, alliances have the incentive to accept militarily powerful members that have a positive effect on their overall military might. This line of reasoning leads to *opposing* expectations about the impact of relative military capabilities on the candidate's and target alliance's incentives. An increase in the capability ratio, which indicates that the candidate state is more powerful compared to the other alliance members, negatively affects a state's incentives to join. This is because the wouldbe member is likely to be a security provider. Conversely, an increase in the capability ratio makes an alliance more likely to accept that state as this would translate into a significant increase in the alliance military might.

Generally, the literature reports a positive relationship between regime type similarity and alliances, although some recent studies have questioned this finding (Siverson & Emmons 1991, Simon & Gartzke 1996, Lai & Reiter 2000). *Joint Democracy* captures the political similarity between the candidate state and target alliance. *Joint Democracy* is calculated as the product between the democratic status of the candidate

respective opportunity set, and then remove the ulterior alliances from the previous membership opportunity sets. For example, if a state forms two alliances in a given year, there would be two alliance membership opportunity sets for that state, but the opportunity set associated with the first alliance formation would not include the subsequent alliance.

¹⁵ The data on states' military capabilities come from Singer, Bremer and Stuckey (1972).

state (0 or 1) and the proportion of member states that are democracies in a given alliance (Poast 2010). A state is coded as a democracy if it scores a six or above on the Polity IV scale for a given year (Marshall, Gurr & Jaggers 2010). Higher values indicate that a democratic state seeks to affiliate with a military alliance that has a large share of democratic members.

As an indicator of geographic proximity, *Geographic Distance* captures both the incentives to ally (in terms of how likely states are to have a common interest or threat) as well as the presumed alliance effectiveness (in terms of how expeditious receiving military support would be). The expectation is that the geographically closer states are, the more likely they are to ally. Measuring geographic distance between states is relatively straightforward in a dyadic setting, and it is usually calculated as the square root distance between the capital cities. This becomes complicated when more than two states are involved. Following Poast (2010), I apply the ‘weakest link’ principle and use the square root distance between the candidate state and the most distant alliance member (Oneal & Russett 1997). For contiguous states, the distance is set to zero (Lai & Reiter 2000, Gibler & Wolford 2006, Poast 2010).¹⁶

Finally, countries’ involvement in international disputes provides an additional indicator about how opportune a given military collaboration is. *Common Threat* captures the idea that countries are more likely to come together if they face a common enemy. *Common Threat* is calculated as the proportion of alliance members that participated in a militarized interstate dispute (MID) against the same third party as the candidate state in the previous ten years (Poast 2010).¹⁷ Empirically, a state’s willingness to join a military alliance should increase with the number of alliance members that fought against the state’s enemies in the past.

In terms of the candidate state’s MID experience, military alliances are interested to know how beligerent the wouldbe member is. One of the costs of accepting new members is the probability that the alliance would be dragged into wars against states with which, absent the new member, the alliance has no quarrels. *Additional Threat* is calculated as the proportion of alliance members that did not participate in a militarized dispute against any of the candidate state’s adversaries in the previous ten years. Discriminating

¹⁶ The data on geographic distance come from Gleditsch and Ward (2001). The data for the period before 1875, as well as the contiguity information come from Gibler and Wolford (2006).

¹⁷ The data on dyadic MIDs come from Maoz (2005).

between the candidate state and target alliance's incentives based on their respective MID records also helps satisfy the MMIwPO's identification requirement. Recall that in order for the model to be identified, it is necessary that the actors' determinant sets are not identical (Poirier 1980, Gordon & Smith 2004).

4.2 Results and Discussion

The results of the empirical analysis are shown in Table 2. The second column presents the estimates from a conditional logit model, the standard univariate logit for multiple choices. These results act as a benchmark against which to compare the results from the MMIwPO. The MMIwPO estimates, which are shown in the third and fourth columns, indicate how covariates influence the choices of the wouldbe member and target alliance separately.

Table 2: The Determinants of Alliance Formation

Regressor	Conditional Logit	Multilateral Mediated Interaction with Partial Observability	
		Wouldbe Member	Target Alliance
Capability Ratio	1.50*** (0.27)	-1.22*** (0.41)	4.06*** (0.59)
Joint Democracy	1.47*** (0.30)	-0.48 (0.41)	5.21*** (0.91)
Geographic Distance	-0.08*** (0.00)	-0.05*** (0.01)	-0.05*** (0.01)
Common Threat	1.19*** (0.26)	0.98*** (0.30)	—
Additional Threat	—	—	-1.18*** (0.34)
Log Likelihood	-1161.99		-1128.74
Number of Observations	11632		11632

Note: The *Conditional Logit* estimates indicate how the regressors influence the likelihood of alliance formation. The *Wouldbe Member* and *Target Alliance* estimates, respectively, indicate how the regressors influence the incentives of the candidate state and target alliance.

The results from the CL model are generally in accord with those from previous studies. Specifically, the likelihood of alliance formation decreases as the candidate state and alliance members are geographically farther apart. This is indicated by the negative and statistically significant coefficient on *Geographic Distance* in the *Conditional Logit* column. The positive and statistically significant coefficients on *Capability Ratio*, *Joint Democracy*, and *Common Threat* indicate that a state is more likely to join a military alliance if that state is militarily strong, the alliance members have a similar regime, or they have previously fought

against the same enemies.


I now turn to the results from the MMIwPO model.¹⁸ The first thing to note is that, unlike the univariate logit, the MMIwPO is able to provide insights into how the covariates affect the choices of *both* the candidate state and the target alliance. It turns out that this has important consequences for the inferences that one would draw. For example, the MMIwPO estimates indicate that a candidate state's military strength has opposing effects on the state's and alliance's incentives. The positive and statistically significant coefficient on *Capability Ratio* in the *Target Alliance* column suggests that an alliance is more likely to accept a militarily strong candidate that has the ability to provide security guarantees to its current members. Contrary, a state is less likely to choose an alliance with militarily weak members. This is because such an alliance is unlikely to provide credible security guarantees. Moreover, as a powerful member of that alliance, the candidate state is likely to be called upon to defend weaker members. Evidence for this line of reasoning comes from the negative and statistically significant coefficient on *Capability Ratio* in the *Wouldbe Member* column. The univariate logit, which reports solely a positive effect, completely misses the negative effect of military capabilities on the probability of alliance formation.

What does it substantively mean that some factors have opposing effects on the candidate state and target alliance? In Figure 2, I use the estimates from the CL and MMIwPO models to show the percentage change in the probability of alliance formation as *Capability Ratio* increases from its mean to one standard deviation above the mean, while all other variables are held at their means. In the case of MMIwPO, I also illustrate the effect of *Capability Ratio* on the state's and alliance's incentives separately. The predicted probabilities are computed from a stylized scenario where a state chooses between two alliances, which in turn decide whether to accept or reject the candidate. The solid vertical lines represent two-tailed 90% confidence intervals, which were computed via simulations based on 10,000 draws from the estimated coefficient vector and variance-covariance matrix.

The CL model suggests that the counterfactual increase in a candidate state's military strength in-

¹⁸ For theoretical reasons, in this analysis I employ the MMIwPO and not the MMIwPO with odds ratio model. When actors' choices are not simultaneous, the disturbances are generally assumed uncorrelated (Abowd & Farber 1982, Przeworski & Vreeland 2002). Arguably, alliance formation is an inherently sequential process. In fact, membership negotiations can take a long time, sometimes even years (Fuller 2014).

Figure 2: The Substantive Effect of Military Capabilities



Note: Figure 2a illustrates the percent change in the probability of alliance formation when *Capability Ratio* increases from its mean to one standard deviation above the mean, while all other variables are held at their means. Figure 2b graphs the substantive effect of *Capability Ratio* on the state's and alliance's incentives separately, as well as the overall effect. The solid square marks indicate the percentage change in probability. The solid vertical lines represent two-tailed 90% confidence intervals.

increases the probability of alliance formation by 22% [15, 26] (Figure 2a). Much more substantive information can be gleaned from the MMIwPO plot. Given the same counterfactual scenario, the MMIwPO estimates indicate that the probability of the target alliance accepting a militarily strong state increases by 48% [39, 56]. Conversely, the probability that a candidate state chooses a weaker alliance decreases by 16% [-25, -9] (Figure 2b). Neither of these substantive quantities of interest can be calculated from the CL estimates. The overall (joint) probability of observing a new alliance increases by 35% [26, 41], 13 percentage points more than with the inappropriate univariate logit. By ignoring the multi-way interaction underpinning an alliance formation scenario, the univariate logit is essentially mixing the effect of military capabilities on the choices of the wouldbe member and target alliance.

As noted above, the CL results indicate that democratic states are more likely to choose alliances with many democratic members. Yet, the MMIwPO estimates suggest that it is the other way around. An alliance with a large share of democratic members is more likely to accept a democratic state, while regime similarity does not play a significant role in a candidate state's decision to join. Evidence for this line of reasoning comes from the positive and statistically significant coefficient on *Joint Democracy* in the *Target Alliance* column, and the statistically insignificant coefficient on the same variable in the *Wouldbe Member*

column. Therefore, the finding of previous studies where joint democracy is a strong indicator of alliance membership, appears to be driven by alliances' ideological considerations, and only to a lesser degree of the candidate state.

Of course it does not have to be the case that all factors that push a candidate state to seek membership in a given alliance make the target alliance less like to accept that state. As expected, the negative and significant coefficients on *Geographic Distance* in the *Wouldbe Member* and *Target Alliance* columns indicate that an increase in the distance between the candidate and alliance members reduces the willingness of both players to cooperate. Lastly, MMIwPO estimates suggest that the MID records of both players affect the probability of alliance formation. In particular, in line with the theoretical expectations, a state is more likely to ally with countries that have the same enemies. This is indicated by the positive and statistically significant coefficient on *Common Threat* in the *Wouldbe Member* column. Conversely, an alliance is less likely to accept a new member that has fought with states that are not enemies of its current members. This is to prevent the alliance being dragged into unwanted disputes. Evidence for this comes from the negative and statistically significant coefficient on *Additional Threat* in the *Target Alliance* column.

5 Conclusions

The conflict in the eastern Ukraine has already claimed more than 30,000 casualties, and left critical civilian infrastructure in ruin (OHCHR 2017). What triggered these events was the Ukrainian executive's abrupt decision to pull back from an already agreed cooperation with the E.U. In the framework of the bilateral interaction between Ukraine and the E.U., it is virtually impossible to explain Ukraine's change of heart. The terms of cooperation, and therefore its value to the signatory parties, did not change in the two months that elapsed from the signing to dismissing the agreement. As a rational actor with consistent preferences, Ukraine should have had the same position on the agreement before and at the Vilnius summit. Reneging the agreement, though, was arguably rational if we acknowledge that for Ukraine the E.U. was not the only alternative. Indeed, Ukraine had been simultaneously negotiating with Russia as well. While the absolute value of the E.U. collaboration did not change, its relative value compared to Russia's counteroffer did. Thus, the conclusion of the Vilnius summit is the result of the multilateral negotiations between Ukraine on one side, and the E.U. and Russia on the other.

Extant empirical estimators are not well equipped to analyze multi-actor interactions. As a result, the common practice is to employ univariate or bivariate models to analyze multilateral phenomena. These approaches are inappropriate, though, for both substantive and methodological reasons. Substantively, scholars are often interested in how the actors engaged in negotiations individually respond to changes in their environment. Univariate models can only provide insights into how factors affect the incentives of a single actor. Studies that employ bivariate estimators can account for the incentives of two actors, but they still ignore the characteristics and motivations of concerned third-parties. Methodologically, analyzing multilateral processes as a set of either monadic or dyadic events biases the results (Croco & Teo 2005, Fordham & Poast 2016, Poast 2010, Signorino 1999).

In this paper, I introduce a new Multilateral Mediated Interaction with Partial Observability model that is specifically designed to analyze multi-actor negotiations. In this model, a player chooses whether to cooperate with any, or none, of several potential partners. For their part, each potential partner decides whether to reject or accept any requests for cooperation. To validate the new estimator, I conduct MC simulations that provide strong evidence for the superior performance of MMIwPO relative to the CL. First, likelihood-ratio tests indicate that employing the more complex model is warranted since MMIwPO fits the data significantly better. Second, MMIwPO consistently outperforms the CL in terms of minimizing the bias of estimated coefficients. Specifically, the CL estimates are always bias, but especially so when a factor simultaneously affects the incentives of more than one actor. Where this is the case, neither the direction nor magnitude of the bias can be inferred from the univariate analysis. Consequently, one cannot adjust the point and interval estimates to account for the bias. Since the new estimator can model the input of multiple actors into a unified, overarching decision making process, the MMIwPO estimates are more consistent.

Lastly, the analysis on alliance formation also highlights the need for a more refined empirical approach. Unlike the CL, MMIwPO is able to distinguish among different types of alliance membership determinants. On the one hand, there are factors that have a similar effect on the candidate state and target alliance (e.g., *Geographic Distance*). On the other hand, some factors have different effects on the actors (e.g., *Joint Democracy*), with some having opposing effects (e.g., *Capability Ratio*). By mixing the determinants' effect on several actors, existing studies essentially report a weighted average of the different effects. They do so by obscuring the exact influence on any one actor, and leading in many cases to incorrect infer-

ences. For example, the CL results suggest that the stronger a state is relative to the alliance members, the more likely it is to seek affiliation with that alliance. MMIwPO estimates, however, suggest that the positive effect from the univariate model is driven by the incentives of target alliances, which are more likely to accept a militarily powerful candidate. In fact, a militarily strong state is less likely to join a weak alliance since it has little to gain from such a commitment.

References

- Abowd, John M. & Henry S. Farber. 1982. "Job Queues and the Union Status of Workers." *Industrial and Labor Relations Review* 35:354–367.
- Agence France-Presse. 2013. "NATO Rules Out Colombia Membership." *Defense News* June 4, 2013. <http://www.defensenews.com/article/20130604/DEFREG02/306040013/NATO-Rules-Out-Colombia-Membership> (Accessed: December 23, 2014).
- Bland, J. Martin & Douglas G. Altman. 2000. "The Odds Ratio." *The BMJ* 320(7247):1468.
- Croco, Sarah & Tze Kwang Teo. 2005. "Assessing the Dyadic Approach to Interstate Conflict Processes: A.k.a. "Dangerous" Dyad-Years." *Conflict Management and Peace Science* 22:5–18.
- Englund, Will & Kathy Lally. 2013. "Ukraine, Under Pressure from Russia, Puts Brakes on E.U. Deal." *The Washington Post* November 21, 2013. http://www.washingtonpost.com/world/europe/ukraine-under-pressure-from-russia-puts-brakes-on-eu-deal/2013/11/21/46c50796-52c9-11e3-9ee6-2580086d8254_story.html (Accessed: January 7, 2015).
- Fordham, Benjamin & Paul Poast. 2016. "All Alliances are Multilateral: Rethinking Alliance Formation in the International System." *Journal of Conflict Resolution* 60(5):840–865.
- Fuller, Liz. 2014. "Georgia Sets Sights On NATO Membership Action Plan." *Radio Free Europe / Radio Liberty* June 16, 2014. <http://www.rferl.org/content/caucasus-report-georgia-nato/25232112.html> (Accessed: December 23, 2014).
- Gibler, Douglas M. & Scott Wolford. 2006. "Alliances, Then Democracy: An Examination of the Relationship between Regime Type and Alliance Formation." *Journal of Conflict Resolution* 50:129–153.
- Gleditsch, Kristian S. & Michael D. Ward. 2001. "Measuring Space: A Minimum Distance Database." *Journal of Peace Research* 38:749–68.
- Gordon, Sanford C. & Alastair Smith. 2004. "Quantitative Leverage Through Qualitative Knowledge: Augmenting the Statistical Analysis of Complex Causes." *Political Analysis* 12:233–255.
- Greene, William. 2003. *Econometric Analysis*. 5th ed. New Jersey: Prentice Hall.
- Heckman, James. 1979. "Sample Selection Bias as a Specification Error." *Econometrica* 47(1):153–161.
- Lai, Brian & Dan Reiter. 2000. "Democracy, Political Similarity, and International Alliances, 1816-1992." *Journal of Conflict Resolution* 44 (2):203–27.
- Leemann, Lucas. 2014. "Strategy and Sample Selection – A Strategic Selection Estimator." *Political Analysis* 22:374–97.
- Lemke, Douglas & William Reed. 2001. "The Relevance of Politically Relevant Dyads." *Journal of Conflict Resolution* 45(1):126–144.
- Logan, John Allen. 1996. "Opportunity and Choice in Socially Structured Labor Markets." *American Journal of Sociology* 102:114–160.

- Logan, John Allen. 1998. "Estimating Two-Sided Logit Models." *Sociological Methodology* 28:139–173.
- Mallén, Patricia Rey. 2013. "Colombia President Santos Approaches NATO; Bolivia, Venezuela And Nicaragua Condemn Him For It." *International Business Times* June 4, 2013. <http://www.ibtimes.com/colombia-president-santos-approaches-nato-bolivia-venezuela-nicaragua-condemn-him-it-1297059> (Accessed: December 23, 2014).
- Maoz, Zeev. 2005. "Dyadic MID Dataset (version 2.0)." <http://vanity.dss.ucdavis.edu/~maoz/dyadmid.html> (Accessed: November 3, 2014).
- Marshall, Monty G., Ted Robert Gurr & Keith Jagers. 2010. "Polity IV Project: Political Regime Characteristics and Transitions, 1800-2009. Dataset Users' Manual." Available from <http://www.systemicpeace.org/inscr/p4manualv2010.pdf>.
- McFadden, Daniel. 1974. Conditional Logit Analysis of Qualitative Choice Behaviour. In *Frontiers in Econometrics*, ed. P. Zarembka. New York: Academic Press pp. 105–42.
- McFadden, Daniel. 1974a. "The Measurement of Urban Travel Demand." *Journal of Public Economics* 3:303–328.
- McLerran, Dale. 2010. "Bivariate Logistic (with Missing Response Values)." *The University of Georgia: SAS-L Archives* June 21, 2010. <https://listserv.uga.edu/cgi-bin/wa?A2=sas-1;7e711936.1006C> (Accessed: January 3, 2015).
- Nieman, Mark David. 2015. "Statistical Analysis of Strategic Interaction with Unobservable Player Actions: Introducing a Strategic Probit with Partial Observability." *Political Analysis* 23(3):429–448.
- Office of the United Nations High Commissioner for Human Rights (OHCHR). 2017. "Report on the human rights situation in Ukraine 16 November 2016 to 15 February 2017." http://www.un.org.ua/images/stories/17th_HRMMU_Report_ENG.pdf (Accessed: September 11, 2017).
- Oneal, John & Bruce Russett. 1997. "The classic liberals were right: Democracy, interdependence, and conflict, 1950-1985." *International Studies Quarterly* 41:267–93.
- Poast, Paul. 2010. "(Mis)Using Dyadic Data to Analyze Multilateral Events." *Political Analysis* 18:403–425.
- Poast, Paul. 2012. "Does Issue Linkage Work? Evidence from European Alliance Negotiations, 1860 to 1945." *International Organization* 66 (2):277–310.
- Poirier, Dale J. 1980. "Partial Observability in Bivariate Probit Models." *Journal of Econometrics* 12:209–217.
- Przeworski, Adam & James Vreeland. 2002. "A Statistical Model of Bilateral Cooperation." *Political Analysis* 10:101–112.
- Signorino, Curtis. 1999. "Strategic Interaction and the Statistical Analysis of International Conflict." *The American Political Science Review* 93:279–97.
- Simon, Michael W. & Erik Gartzke. 1996. "Political System Similarity and the Choice of Allies: Do Democracies Flock Together or Do Opposites Attract?" *Journal of Conflict Resolution* 40 (4):617–35.

- Singer, J. David, Stuart Bremer & John Stuckey. 1972. Capability Distribution, Uncertainty, and Major Power War, 1820-1965. In *Peace, War, and Numbers*, ed. Bruce Russett. Beverly Hills: Sage pp. 19–48.
- Siverson, Randolph M. & Juliann Emmons. 1991. “Birds of a Feather: Democratic Political Systems and Alliance Choices in the Twentieth Century.” *Journal of Conflict Resolution* 35 (2):285–306.
- Szumilas, Magdalena. 2010. “Explaining Odds Ratios.” *Journal of the Canadian Academy of Child and Adolescent* 19(3):227–229.
- Train, Kenneth E. 2007. *Discrete Choice Models with Simulation*. New York: Cambridge University Press.
- Wooldridge, Jeffrey M. 2003. “Cluster-Sample Methods in Applied Econometrics.” *American Economic Review* 93:133–138.

Online Appendix A: Deriving the MMIwPO Likelihood Function

In Section 2.1 of the main text I note that the simplified conditional probability that actor i cooperates with actor j , and not with $\sim j$, out of a set of K_i potential partners is

$$\Pr(y_{ij} = 1 | x_{i1}, \dots, x_{iK_i}, z_{1i}, \dots, z_{K_i i}; \sum_{k=1}^{K_i} y_{ik} = 1) = \frac{\frac{e^{x_{ij}\beta} e^{z_{ji}\gamma}}{1 + e^{x_{ij}\beta} + e^{z_{ji}\gamma}}}{\sum_{k=1}^{K_i} \frac{e^{x_{ik}\beta} e^{z_{ki}\gamma}}{1 + e^{x_{ik}\beta} + e^{z_{ki}\gamma}}} \quad (\text{A1})$$

where $j \in K_i$.

I now present step-by-step how the conditional probability underlying the MMIwPO's likelihood function is computed. Recall that the joint probability that actor i seeks to cooperate with j *and* that actor j accepts to cooperate with i is

$$\Pr(o_{ij} | x_{ij}, z_{ji} = 1) = P_{ij} \times P_{ji} \quad (\text{A2})$$

where $P_{ij} = F(x_{ij}\beta)$ and $P_{ji} = F(z_{ji}\gamma)$. P_{ij} represents the individual probability that i wants to cooperate with j , β is a vector of coefficients, and x_{ij} is a vector of independent variables for actor i associated with potential partner j . P_{ji} represents the individual probability that j wants to cooperate with i , γ is a vector of coefficients, and z_{ji} is a vector of independent variables for actor j associated with potential partner i . $F(\cdot)$ is the cumulative distribution function of the standard logistic distribution, i.e., $F(\cdot) = \Lambda(\cdot)$.

Conversely, the probability that actor i does not cooperate with actor j , either because i refuses to cooperate with j or because j rejects i (or both), is simply

$$\Pr(o_{ij} | x_{ij}, z_{ji} = 0) = 1 - (P_{ij} \times P_{ji}). \quad (\text{A3})$$

Starting from the joint dyadic probabilities, we can write the *conditional* probability that actor i

cooperates with actor j , and not with $\sim j$, out of a set of K_i potential partners as

$$\Pr(y_{ij} = 1, y_{i2} = 0, \dots, y_{iK_i} = 0 \mid y_{ij} + y_{i2} + \dots + y_{iK_i} = 1) = \frac{\Pr(y = (1, 0, \dots, 0)')}{\Pr(y = (1, 0, \dots, 0)') + \Pr(y = (0, 1, \dots, 0)') + \dots + \Pr(y = (0, 0, \dots, 1)')} \quad (\text{A4})$$

where $j \in K_i$.

Using the joint probabilities in Eq. (A2) and Eq. (A3), the numerator in Eq. (A4) can be computed as

$$\begin{aligned} & \left\{ [P_{ij}P_{ji}] \times [1 - P_{i2}P_{2i}] \times \dots \times [1 - P_{iK_i}P_{K_i i}] \right\} = \\ & \left\{ \left[\left(\frac{e^{x_{ij}\beta}}{1 + e^{x_{ij}\beta}} \right) \left(\frac{e^{z_{ji}\gamma}}{1 + e^{z_{ji}\gamma}} \right) \right] \times \left[1 - \left(\frac{e^{x_{i2}\beta}}{1 + e^{x_{i2}\beta}} \right) \left(\frac{e^{z_{2i}\gamma}}{1 + e^{z_{2i}\gamma}} \right) \right] \times \dots \times \left[1 - \left(\frac{e^{x_{iK_i}\beta}}{1 + e^{x_{iK_i}\beta}} \right) \left(\frac{e^{z_{K_i i}\gamma}}{1 + e^{z_{K_i i}\gamma}} \right) \right] \right\} = \\ & \left\{ \left[\frac{e^{x_{ij}\beta} e^{z_{ji}\gamma}}{(1 + e^{x_{ij}\beta})(1 + e^{z_{ji}\gamma})} \right] \times \left[\frac{1 + e^{x_{i2}\beta} + e^{z_{2i}\gamma}}{(1 + e^{x_{i2}\beta})(1 + e^{z_{2i}\gamma})} \right] \times \dots \times \left[\frac{1 + e^{x_{iK_i}\beta} + e^{z_{K_i i}\gamma}}{(1 + e^{x_{iK_i}\beta})(1 + e^{z_{K_i i}\gamma})} \right] \right\} = \\ & \left\{ \frac{e^{x_{ij}\beta} e^{z_{ji}\gamma} \times (1 + e^{x_{i2}\beta} + e^{z_{2i}\gamma}) \times \dots \times (1 + e^{x_{iK_i}\beta} + e^{z_{K_i i}\gamma})}{(1 + e^{x_{ij}\beta})(1 + e^{z_{ji}\gamma}) \times (1 + e^{x_{i2}\beta})(1 + e^{z_{2i}\gamma}) \times \dots \times (1 + e^{x_{iK_i}\beta})(1 + e^{z_{K_i i}\gamma})} \right\}, \quad (\text{A5}) \end{aligned}$$

and the denominator as

$$\begin{aligned} & \left\{ [P_{ij}P_{ji}] \times [1 - P_{i2}P_{2i}] \times \dots \times [1 - P_{iK_i}P_{K_i i}] \right\} + \left\{ [1 - P_{ij}P_{ji}] \times [P_{i2}P_{2i}] \times \dots \times [1 - P_{iK_i}P_{K_i i}] \right\} + \\ & \dots + \left\{ [1 - P_{ij}P_{ji}] \times [1 - P_{i2}P_{2i}] \times \dots \times [P_{iK_i}P_{K_i i}] \right\} = \\ & \left\{ \left[\left(\frac{e^{x_{ij}\beta}}{1 + e^{x_{ij}\beta}} \right) \left(\frac{e^{z_{ji}\gamma}}{1 + e^{z_{ji}\gamma}} \right) \right] \times \left[1 - \left(\frac{e^{x_{i2}\beta}}{1 + e^{x_{i2}\beta}} \right) \left(\frac{e^{z_{2i}\gamma}}{1 + e^{z_{2i}\gamma}} \right) \right] \times \dots \times \left[1 - \left(\frac{e^{x_{iK_i}\beta}}{1 + e^{x_{iK_i}\beta}} \right) \left(\frac{e^{z_{K_i i}\gamma}}{1 + e^{z_{K_i i}\gamma}} \right) \right] \right\} + \\ & \left\{ \left[1 - \left(\frac{e^{x_{ij}\beta}}{1 + e^{x_{ij}\beta}} \right) \left(\frac{e^{z_{ji}\gamma}}{1 + e^{z_{ji}\gamma}} \right) \right] \times \left[\left(\frac{e^{x_{i2}\beta}}{1 + e^{x_{i2}\beta}} \right) \left(\frac{e^{z_{2i}\gamma}}{1 + e^{z_{2i}\gamma}} \right) \right] \times \dots \times \left[1 - \left(\frac{e^{x_{iK_i}\beta}}{1 + e^{x_{iK_i}\beta}} \right) \left(\frac{e^{z_{K_i i}\gamma}}{1 + e^{z_{K_i i}\gamma}} \right) \right] \right\} + \\ & \dots + \\ & \left\{ \left[1 - \left(\frac{e^{x_{ij}\beta}}{1 + e^{x_{ij}\beta}} \right) \left(\frac{e^{z_{ji}\gamma}}{1 + e^{z_{ji}\gamma}} \right) \right] \times \left[1 - \left(\frac{e^{x_{i2}\beta}}{1 + e^{x_{i2}\beta}} \right) \left(\frac{e^{z_{2i}\gamma}}{1 + e^{z_{2i}\gamma}} \right) \right] \times \dots \times \left[\left(\frac{e^{x_{iK_i}\beta}}{1 + e^{x_{iK_i}\beta}} \right) \left(\frac{e^{z_{K_i i}\gamma}}{1 + e^{z_{K_i i}\gamma}} \right) \right] \right\} = \\ & \left\{ \left[\frac{e^{x_{ij}\beta} e^{z_{ji}\gamma}}{(1 + e^{x_{ij}\beta})(1 + e^{z_{ji}\gamma})} \right] \times \left[\frac{1 + e^{x_{i2}\beta} + e^{z_{2i}\gamma}}{(1 + e^{x_{i2}\beta})(1 + e^{z_{2i}\gamma})} \right] \times \dots \times \left[\frac{1 + e^{x_{iK_i}\beta} + e^{z_{K_i i}\gamma}}{(1 + e^{x_{iK_i}\beta})(1 + e^{z_{K_i i}\gamma})} \right] \right\} + \end{aligned}$$

$$\begin{aligned}
& \left\{ \left[\frac{1 + e^{x_{ij}\beta} + e^{z_{ji}\gamma}}{(1 + e^{x_{ij}\beta})(1 + e^{z_{ji}\gamma})} \right] \times \left[\frac{e^{x_{i2}\beta} e^{z_{2i}\gamma}}{(1 + e^{x_{i2}\beta})(1 + e^{z_{2i}\gamma})} \right] \times \dots \times \left[\frac{1 + e^{x_{iK_i}\beta} + e^{z_{K_i i}\gamma}}{(1 + e^{x_{iK_i}\beta})(1 + e^{z_{K_i i}\gamma})} \right] \right\} + \\
& \quad , \dots , + \\
& \left\{ \left[\frac{1 + e^{x_{ij}\beta} + e^{z_{ji}\gamma}}{(1 + e^{x_{ij}\beta})(1 + e^{z_{ji}\gamma})} \right] \times \left[\frac{1 + e^{x_{i2}\beta} + e^{z_{2i}\gamma}}{(1 + e^{x_{i2}\beta})(1 + e^{z_{2i}\gamma})} \right] \times \dots \times \left[\frac{e^{x_{iK_i}\beta} e^{z_{K_i i}\gamma}}{(1 + e^{x_{iK_i}\beta})(1 + e^{z_{K_i i}\gamma})} \right] \right\} = \\
& \quad \left\{ \frac{e^{x_{ij}\beta} e^{z_{ji}\gamma} \times (1 + e^{x_{i2}\beta} + e^{z_{2i}\gamma}) \times \dots \times (1 + e^{x_{iK_i}\beta} + e^{z_{K_i i}\gamma})}{(1 + e^{x_{ij}\beta})(1 + e^{z_{ji}\gamma}) \times (1 + e^{x_{i2}\beta})(1 + e^{z_{2i}\gamma}) \times \dots \times (1 + e^{x_{iK_i}\beta})(1 + e^{z_{K_i i}\gamma})} \right\} + \\
& \quad \left\{ \frac{(1 + e^{x_{ij}\beta} e^{z_{ji}\gamma}) \times e^{x_{i2}\beta} e^{z_{2i}\gamma} \times \dots \times (1 + e^{x_{iK_i}\beta} + e^{z_{K_i i}\gamma})}{(1 + e^{x_{ij}\beta})(1 + e^{z_{ji}\gamma}) \times (1 + e^{x_{i2}\beta})(1 + e^{z_{2i}\gamma}) \times \dots \times (1 + e^{x_{iK_i}\beta})(1 + e^{z_{K_i i}\gamma})} \right\} + \\
& \quad , \dots , + \\
& \quad \left\{ \frac{(1 + e^{x_{ij}\beta} + e^{z_{ji}\gamma}) \times (1 + e^{x_{i2}\beta} + e^{z_{2i}\gamma}) \times \dots \times e^{x_{iK_i}\beta} e^{z_{K_i i}\gamma}}{(1 + e^{x_{ij}\beta})(1 + e^{z_{ji}\gamma}) \times (1 + e^{x_{i2}\beta})(1 + e^{z_{2i}\gamma}) \times \dots \times (1 + e^{x_{iK_i}\beta})(1 + e^{z_{K_i i}\gamma})} \right\}. \tag{A6}
\end{aligned}$$

If we multiply the numerator, Eq. (A5), and denominator, Eq. (A6), by

$$\frac{(1 + e^{x_{ij}\beta})(1 + e^{z_{ji}\gamma}) \times (1 + e^{x_{i2}\beta})(1 + e^{z_{2i}\gamma}) \times \dots \times (1 + e^{x_{iK_i}\beta})(1 + e^{z_{K_i i}\gamma})}{(1 + e^{x_{ij}\beta} + e^{z_{ji}\gamma}) \times (1 + e^{x_{i2}\beta} + e^{z_{2i}\gamma}) \times \dots \times (1 + e^{x_{iK_i}\beta} + e^{z_{K_i i}\gamma})},$$

the conditional probability can be written as

$$\frac{\frac{e^{x_{ij}\beta} e^{z_{ji}\gamma}}{1 + e^{x_{ij}\beta} + e^{z_{ji}\gamma}}}{\frac{e^{x_{ij}\beta} e^{z_{ji}\gamma}}{1 + e^{x_{ij}\beta} + e^{z_{ji}\gamma}} + \frac{e^{x_{i2}\beta} e^{z_{2i}\gamma}}{1 + e^{x_{i2}\beta} + e^{z_{2i}\gamma}} + \dots + \frac{e^{x_{iK_i}\beta} e^{z_{K_i i}\gamma}}{1 + e^{x_{iK_i}\beta} + e^{z_{K_i i}\gamma}}} = \frac{\frac{e^{x_{ij}\beta} e^{z_{ji}\gamma}}{1 + e^{x_{ij}\beta} + e^{z_{ji}\gamma}}}{\sum_{k=1}^{K_i} \frac{e^{x_{ik}\beta} e^{z_{ki}\gamma}}{1 + e^{x_{ik}\beta} + e^{z_{ki}\gamma}}} \tag{A7}$$

where $j \in K_i$.

Online Appendix B: Deriving the Odds Ratio for a Bivariate Logit

In Section 2.2 of the main text, when computing the joint outcome P_{oij} as a function of the odds ratio, ψ , and the marginal individual probabilities, P_{ij} and P_{ji} , I jump straight to the solution of the quadratic equation. Here I present the intermediate steps that link the formula for computing the odds ratio, ψ , to the quadratic equation in P_{oij} . Recall that the odds ratio is defined as $\psi = \frac{\Pr(P_{ij}=1 | P_{ji}=1) \times \Pr(P_{ij}=0 | P_{ji}=0)}{\Pr(P_{ij}=1 | P_{ji}=0) \times \Pr(P_{ij}=0 | P_{ji}=1)}$, and that $P_{oij} = \Pr(P_{ij} = 1 | P_{ji} = 1) = P_{ij} \times P_{ji}$.

$$\begin{aligned}
 \psi &= \frac{\Pr(P_{ij} = 1 | P_{ji} = 1) \times \Pr(P_{ij} = 0 | P_{ji} = 0)}{\Pr(P_{ij} = 1 | P_{ji} = 0) \times \Pr(P_{ij} = 0 | P_{ji} = 1)} \\
 \psi &= \frac{P_{oij} \times (1 - P_{ij} - P_{ji} + P_{oij})}{(P_{ij} - P_{oij}) \times (P_{ji} - P_{oij})} \\
 0 &= P_{oij} \times (1 - P_{ij} - P_{ji} + P_{oij}) - \psi \times (P_{ij} - P_{oij}) \times (P_{ji} - P_{oij}) \\
 0 &= P_{oij} - P_{oij} \times P_{ij} - P_{oij} \times P_{ji} + P_{oij}^2 - \psi \times (P_{ij} \times P_{ji} - P_{oij} \times P_{ij} - P_{oij} \times P_{ji} + P_{oij}^2) \\
 0 &= P_{oij} - P_{oij} \times (P_{ij} + P_{ji}) + P_{oij}^2 - \psi \times [P_{ij} \times P_{ji} - P_{oij} \times (P_{ij} + P_{ji}) + P_{oij}^2] \\
 0 &= P_{oij} - P_{oij} \times (P_{ij} + P_{ji}) + P_{oij}^2 - \psi \times P_{ij} \times P_{ji} + \psi \times P_{oij} \times (P_{ij} + P_{ji}) - \psi \times P_{oij}^2 \\
 0 &= P_{oij}^2 - \psi \times P_{oij}^2 + P_{oij} - P_{oij} \times (P_{ij} + P_{ji}) + \psi \times P_{oij} \times (P_{ij} + P_{ji}) - \psi \times P_{ij} \times P_{ji} \\
 0 &= (1 - \psi) \times P_{oij}^2 + [1 - (P_{ij} + P_{ji}) + \psi \times (P_{ij} + P_{ji})] \times P_{oij} - \psi \times P_{ij} \times P_{ji} \\
 0 &= (1 - \psi) \times P_{oij}^2 + [1 - (1 - \psi) \times (P_{ij} + P_{ji})] \times P_{oij} - \psi \times P_{ij} \times P_{ji} \\
 0 &= a \times P_{oij}^2 + b \times P_{oij} + c.
 \end{aligned} \tag{B8}$$

The solution to the quadratic equation is

$$P_{oij} = \left(-b + \sqrt{b^2 - 4 \times a \times c} \right) / (2 \times a) \tag{B9}$$

where $a = 1 - \psi$, $\psi \neq 1$, $b = 1 - (1 - \psi) \times (P_{ij} + P_{ji})$, and $c = -\psi \times P_{ij} \times P_{ji}$.

Online Appendix C: The Full Set of Monte Carlo Simulations

In Section 3 of the main text, I discuss only a subset of the Monte Carlo (MC) simulations that I conduct to validate the new empirical estimator. I now briefly review the set-up of the simulations, and then present the results from all MC experiments. In each MC experiment, 2,000 players choose whether they want to cooperate with any, or none, of 5 potential partners. Similarly, each potential partner decides whether to reject or accept a player were he willing to cooperate. Both the players and potential partners choose the option that maximizes the utility accrued from that cooperation, or the lack thereof. Since there are two sets of actors (i.e., the players (N) and potential partners (K)), there are two distinct utility functions:

$$y_{ij}^* = \beta_1 IV1_{ij} + \beta_2 IV2_{ij} + \epsilon_{ij} \quad (C10)$$

$$y_{ji}^* = \gamma_1 IV1_{ji} + \gamma_2 IV3_{ji} + \epsilon_{ji} \quad (C11)$$

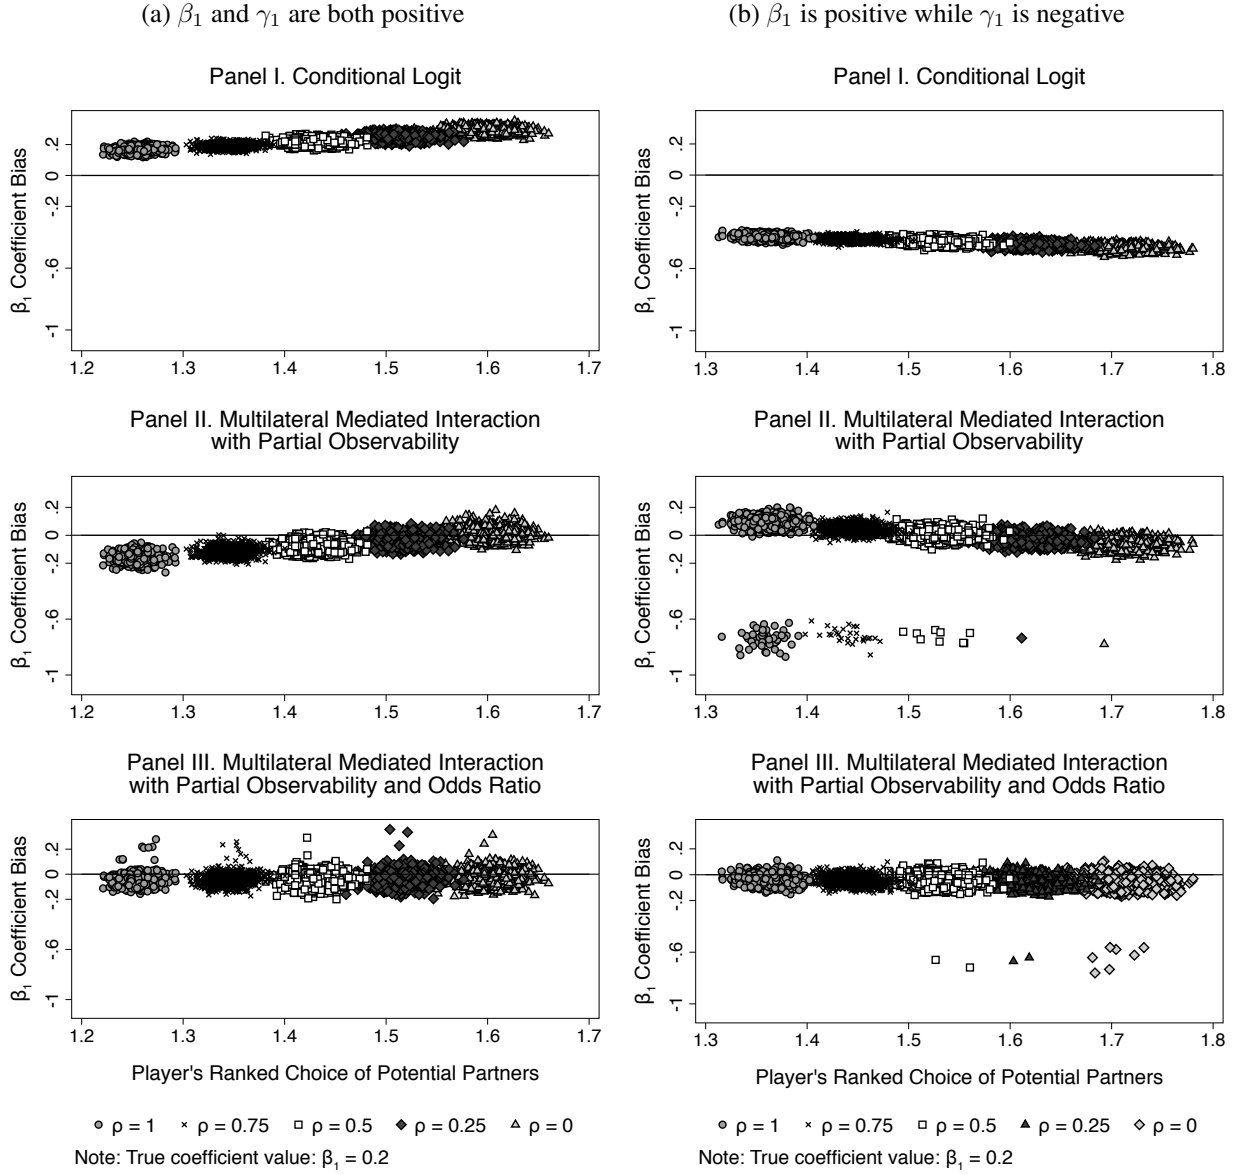
where $i \in N, j \in K, \beta_1 = 0.2, \beta_2 = -0.6, \gamma_1 = +/ - 0.8, \gamma_2 = 0.4$, and $\text{cov}(\epsilon_{ij}, \epsilon_{ji}) = \rho$. All three independent variables have values randomly drawn from the standard logistic distribution, with a mean of 0 and variance of $\frac{\pi^2}{3}$. The error terms are randomly drawn from a cumulative logistic distribution with a correlation coefficient of ρ .

One assumption of the MMIwPO model is that potential partners are substitutes not complements. In practical terms this means that actor i can choose and be accepted by only one alternative from the choice set K_i . Therefore, the dependent variable is coded one for the (y_{ij}^*, y_{ji}^*) dyad that maximizes the product of the two latent variables, given that both individual probabilities equal one. Formally, the dependent variable, y_{ij} , equals one if $\left((y_{ij}^* \times y_{ji}^*) = \max_{j \in K_i} \left\{ (y_{i1}^* \times y_{1i}^*), \dots, (y_{iK_i}^* \times y_{K_i i}^*) \right\} \mid (y_{ij} = 1 \mid y_{ji} = 1) \right)$, zero otherwise.

Overall, there are thirty MC experiments, which represent the combination between (i) five different values for the disturbance correlation parameter (i.e., ρ successively takes the value of 0, 0.25, 0.5, 0.75, and 1), (ii) two different values for the γ_1 coefficient ($+/- 0.8$), and (iii) three different estimators (CL, MMIwPO, MMIwPO with Odds Ratio). For each of the thirty scenarios, I run 1,000 simulations.

I employ two indicators to discriminate between the competing estimators. First, I perform a likelihood-ratio test to assess the goodness-of-fit between the CL and the two MMIwPO models. With no exception,

Figure C1: The β_1 Coefficient Bias



Note: Figure C1 illustrates the bias in the estimated β_1 coefficient for three different models: Conditional Logit, Multilateral Mediated Interaction with Partial Observability, and Multilateral Mediated Interaction with Partial Observability and Odds Ratio. The left panels indicate scenarios where the common factor, $IV1$, has a positive effect on both sets of actors, N and K . The right panels indicate scenarios where $IV1$ has a positive effect on the N actors, and a negative effect on the K s. The estimates are based on 1,000 MC simulations.

both MMIwPO and MMIwPO with Odds Ratio fit the simulated data significantly better in every MC experiment. This test indicates that the increased complexity of the new Multilateral Mediated Interaction with Partial Observability model is justified in terms of the significant improvement in fit.

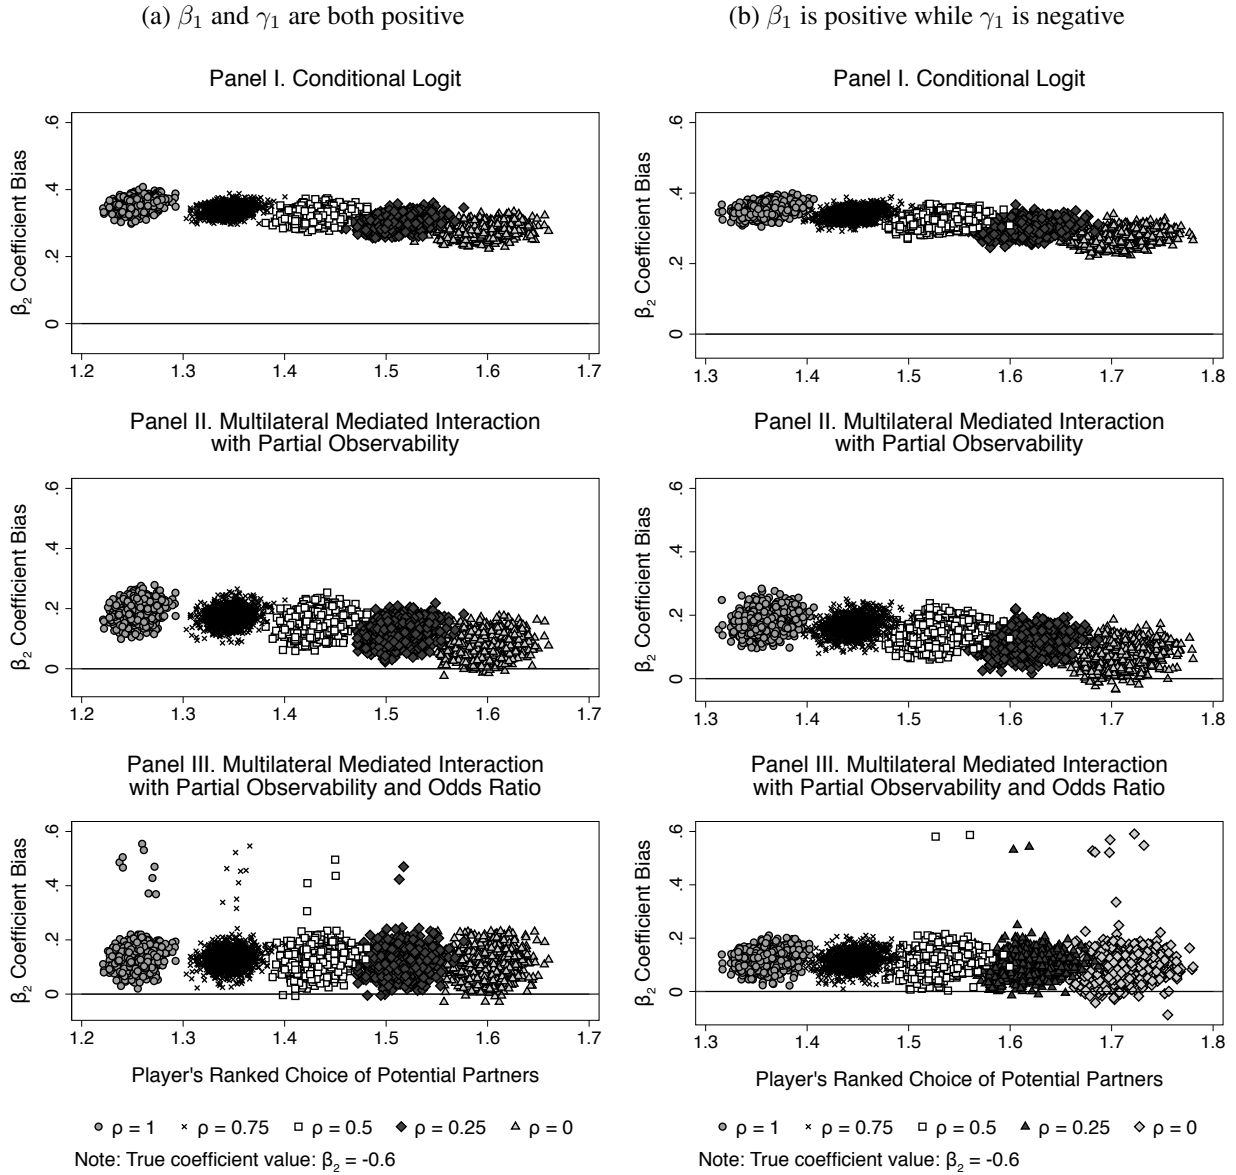
Second, I contrast the three models in terms of their ability to minimize the coefficients' bias. Fig-

ure C1 illustrates the bias in the β_1 coefficient for each of the three estimators across the player's ranked choice of potential partners, at five levels of correlation between disturbances. The coefficient bias is simply the difference between the estimated and the true value of that coefficient. The ranked choice captures potential partners' position on the player's preference scale. For example, a value of one indicates that the player was accepted by the partner that he ranks first. Higher values indicate that the player was rejected by his preferred choice, and he had to settle for his second best alternative, or worse. The left panels indicate scenarios where the common factor, $IV1$, has a positive effect on both sets of actors, N and K . The right panels indicate scenarios where $IV1$ has a positive effect on the N actors, and a negative effect on the K s.

Across all scenarios, the MMIwPO with Odds Ratio reports the least biased coefficients, with the bias values straddling the zero line (Figure C1 Panel III). The next best performing model is the MMIwPO. Since this model does not account for the interdependence between actors' choices, it performs best when the correlation is relatively low. As the disturbance correlation level increases, the estimated coefficients become increasingly biased (Figure C1 Panel II). Lastly, the CL estimates are always biased regardless of the disturbance correlation level. To make things worse, the direction of the β_1 bias is determined by the sign of the γ_1 coefficient from the potential partners' utility function. Recall that β_1 is always positive. When γ_1 is also positive, the estimated β_1 coefficient is always bigger than its true value since it is artificially pulled up by the positive effect that $IV1$ has on the potential partners' incentives (Figure C1a Panel I). When γ_1 is negative, the CL's β_1 coefficient is always smaller than its true value since it is pulled down by $IV1$'s negative effect on the potential partners' incentives (Figure C1b Panel I). In effect, the conditional logit model reports a weighted average of the $IV1$'s effect on all actors.

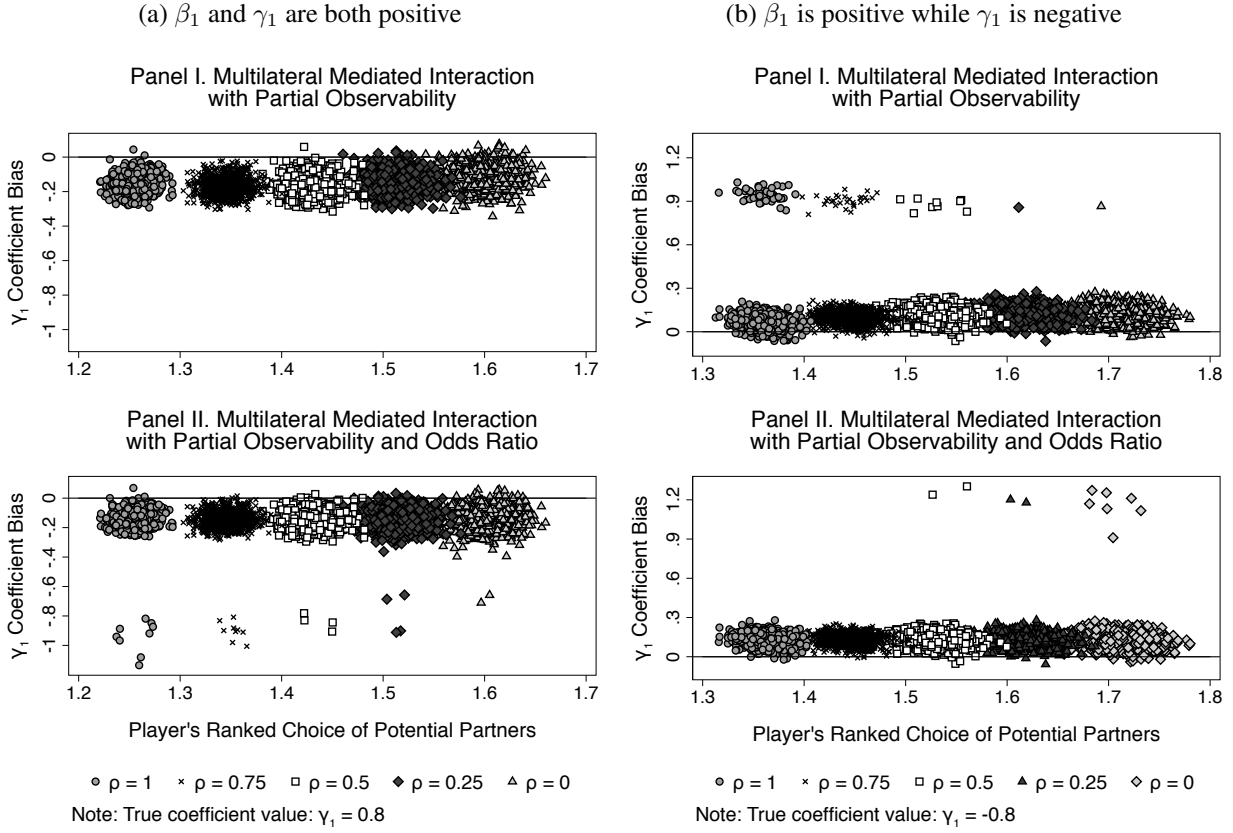
Figure C2 illustrates the β_2 bias, the coefficient on $IV2$, a variable that affects only players' incentives. Not surprisingly, the estimated coefficients are more consistent across all models. This indicates that, in contrast to the β_1 estimate, β_2 is not directly affected by the fluctuation induced by shifting γ_1 's sign. Overall, the insights are similar to the ones from the β_1 coefficient analysis. The MMIwPO with Odds Ratio is the model that reports the least biased coefficients, followed by MMIwPO, and then the CL. For the last two models, the estimated coefficients are increasingly biased as the correlation level increases.

Figure C2: The β_2 Coefficient Bias



Note: Figure C2 illustrates the bias in the estimated β_2 coefficient for three different models: Conditional Logit, Multilateral Mediated Interaction with Partial Observability, and Multilateral Mediated Interaction with Partial Observability and Odds Ratio. The left panels indicate scenarios where the common factor, $IV1$, has a positive effect on both sets of actors, N and K . The right panels indicate scenarios where $IV1$ has a positive effect on the N actors, and a negative effect on the K s. The estimates are based on 1,000 MC simulations.

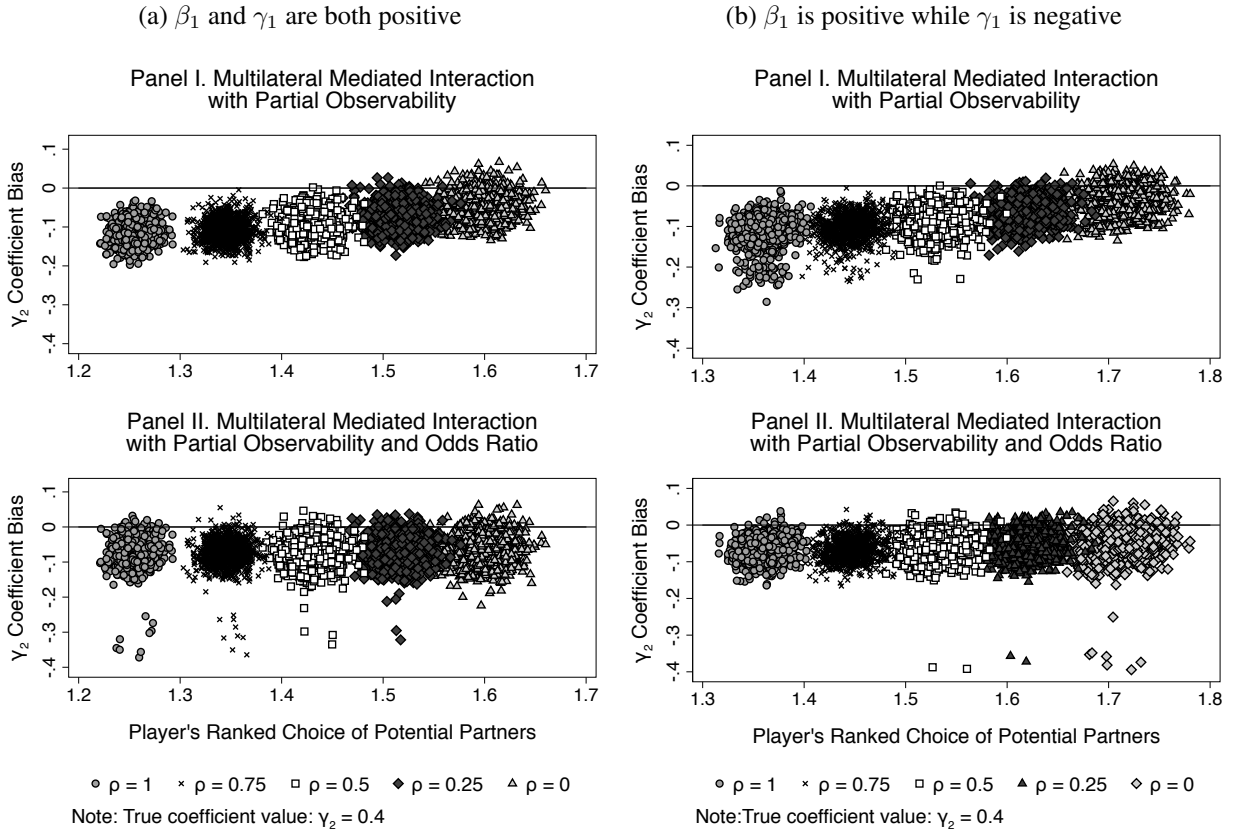
Figure C3: The γ_1 Coefficient Bias



Note: Figure C3 illustrates the bias in the estimated γ_1 coefficient for two models: Multilateral Mediated Interaction with Partial Observability, and Multilateral Mediated Interaction with Partial Observability and Odds Ratio. The left panels indicate scenarios where the common factor, IV_1 , has a positive effect on both sets of actors, N and K . The right panels indicate scenarios where IV_1 has a positive effect on the N actors, and a negative effect on the K s. The estimates are based on 1,000 MC simulations.

Finally, Figure C3 and C4 illustrate the bias in the γ_1 and γ_2 coefficients, respectively. The first thing to note is that there are no coefficient bias estimates for the CL model. As all univariate models, the conditional logit reports only one set of coefficients, and, consequently, we cannot estimate the covariates' effect on individual actors. This means that, when it comes to potential partners' determinants, we can only compare the MMIwPO and MMIwPO with Odds Ratio models. Figure C3 and C4 indicate that, while on average both models perform fairly well, where actors' decisions are interrelated, the MMIwPO with Odds Ratio's coefficients are less biased. In the next section, I discuss in greater detail the tradeoffs between using one or the other version of the Multilateral Mediated Interaction with Partial Observability model.

Figure C4: The γ_2 Coefficient Bias



Note: Figure C4 illustrates the bias in the estimated γ_2 coefficient for two models: Multilateral Mediated Interaction with Partial Observability, and Multilateral Mediated Interaction with Partial Observability and Odds Ratio. The left panels indicate scenarios where the common factor, $IV1$, has a positive effect on both sets of actors, N and K . The right panels indicate scenarios where $IV1$ has a positive effect on the N actors, and a negative effect on the K s. The estimates are based on 1,000 MC simulations.

Online Appendix D: MMIwPO vs. MMIwPO with Odds Ratio Model

When choosing whether to employ the MMIwPO or the MMIwPO with Odds Ratio model, one ought to consider both theoretical and practical matters. From a theoretical perspective, the relevant question is whether the researcher has any priors or theoretical convictions with respect to the sequence in which actors make their decisions. Generally, there are two alternative approaches to estimating multivariate models. One approach assumes that actors' choices occur sequentially and that the disturbances are uncorrelated (Abowd & Farber 1982, Przeworski & Vreeland 2002). A second approach allows actors' choices to (potentially) occur simultaneously, in which case the disturbances are correlated (Poirier 1980). Thus, theoretically, one should employ the MMIwPO when analyzing cases that correspond to the first scenario, and the MMIwPO with Odds Ratio where cases fit the latter scenario.

If the researcher is agnostic with respect to actors' sequence of moves, a safe approach is to first run the MMIwPO with Odds Ratio. The nice feature of this model is that it allows one to empirically test the disturbance interdependence assumption. If the odds ratio parameter is statistically insignificant, we can employ the simple MMIwPO specification. If the estimate is statistically significant, though, one should employ the more complex estimation procedure.

There are, however, a couple of practical problems associated with the MMIwPO with Odds Ratio. One drawback is that this model specification has a lower convergence rate. In the MC experiments, the MMIwPO has a 100% convergence rate. The overall convergence rate for the MMIwPO with Odds Ratio model is still fairly high, but it varies across scenarios. Specifically, when the common independent variable has a similar effect on the actors, the average convergence rate is 99.84% (min = 99.3% , when $\rho = 1$; max = 100%, when $\rho = 0$). Where *IV1* has different effects on the actors, the average convergence rate is 87.62% (min = 82% , when $\rho = 1$; max = 93.9%, when $\rho = 0$). To some extent, these average convergence rates may be overly optimistic. In the MC experiments there are only five potential partners, four covariates, and the number of observations is capped at 10,000. A substantial increase in any of these parameters is likely to lead to lower convergence rates.

The second cost of employing the MMIwPO with Odds Ratio model is the increased computational time. Take, for example, the MC simulation with the highest number of observations, 9,075.¹ When estimat-

¹ The exact specifications of this MC iteration are: 1,815 players, five potential partners, four independent

ing this scenario, the MMIwPO reached convergence in two seconds, while the MMIwPO with Odds Ratio after fifty-three seconds. In fact, the computational time for the MMIwPO with Odds Ratio model increases exponentially with the number of independent variables and that of available alternatives. Of course, the more elaborate a model is, the higher the risk of not reaching convergence at all.

variables (three distinct), and a disturbance correlation level of one.