

Measuring Latent Ties from Rich but Noisy Observational Data

Yinxuan Wang* Bokai Qi†

Version: October 6, 2026

Abstract

Political scientists often rely on proxy networks to study relationships such as patronage, trust, and coalitional alignment that are not directly observed. Yet proxies based on co-working, shared hometowns, schools, or clubs may capture only part of these relationships and be unevenly observed across dyads. We propose *Multi-indicator Latent Tie Inference* (MuLTI), which combines relational traces and dyadic information to estimate distinct latent ties. We establish nonparametric identification using conditionally independent, tie-specific anchors and estimate the model by maximum likelihood with stratified stochastic optimization. Simulations show that MuLTI outperforms proxy-based benchmarks, especially when indicators are noisy and multiple indicators are available. Applied to 4,424 Chinese Communist revolutionary elites and more than 9.7 million dyads, MuLTI reveals a positive association between wartime adversarial exposure and later Cultural Revolution mortality. MuLTI provides a scalable framework for estimating distinct latent relationships from imperfect sources of information that conventional proxy coding cannot capture.

[Word Count: 9709]

Please do not circulate without authors' written permission.

*Ph.D. candidate, Wilf Family Department of Politics, New York University. Email: yinxuan.wang@nyu.edu.

†Ph.D. candidate, Wilf Family Department of Politics, New York University. Email: bokai.qi@nyu.edu.

We thank Cassandra Handan-Nader, Arturas Rozenas, Cyrus Samii, Tara Slough, Luca Vitale, Munroe Kim, Lena Mangold, Jingyuan Qian, Tony Yang, and participants at NYU Methods Workshop, NYU Comparative Workshop, MPSA 2026, PolMeth 2026, and APSA 2026 for the valuable suggestions and feedback. All errors remain our own.

1 Introduction

Political scientists increasingly use social network analysis to study political phenomena (Wood, 2008; Lazer, 2011; Ward, Stovel and Sacks, 2011; Hall and Lamont, 2013; Campbell, 2013; Cranmer et al., 2017; Cruz, Labonne and Querubin, 2020; Larson, 2021; Zaytsev et al., 2023), but the ties that make networks theoretically meaningful are often not directly observed. Studies of elite interaction, repression, clientelism, protest, public goods provision, civil conflict, and state building frequently depend on relational concepts such as trust and reciprocity (Abramson, Carter and Ying, 2022; Masterson, 2024), patronage and clientelism (Hicken, 2011; Shami, 2012; Jiang, 2018; Ravanilla, Sexton and Haim, 2022), and monitoring (Brierley and Nathan, 2021). We refer to these unobserved relationships of theoretical interest as *latent ties*.

In some research settings, network data can be elicited directly. In surveys, field and lab experiments, researchers may ask respondents to name people with whom they have particular relationships and, in some designs, to report relationships among those people (Larson and Lewis, 2020). But many political networks cannot be measured in this way. Researchers cannot ask historical actors, such as authoritarian elites, rebels, bureaucrats, party cadres, or members of clandestine organizations, whom they trusted, feared, supported, monitored, or followed. As Table 1 summarizes, empirical political network research therefore faces different measurement environments: some data sources permit direct elicitation of networks, some provide direct relational records, but many observational sources provide only indirect traces from which *latent ties* must be inferred.

When researchers cannot directly elicit ties, a common strategy is to construct a proxy network from one or more readily observable relations.¹ Many studies use homophilous attributes as proxies for social ties with an intuitive and often useful premise: similarity breeds connection (McPherson, Smith-Lovin and Cook, 2001). If two nodes share the same hometown, school, club, work unit, patron, or even family name, then they had greater opportunity and perhaps stronger propensity to form a tie. This logic underlies a large body of work on elite politics, bureaucratic promotion,

¹As Perliger and Pedahzur (2011) summarized, these observable relations can include acquaintances, close friendships, family ties, and time spent together.

Table 1: Relational Information in Empirical Network Data

Setting	Observed	Examples
Elicited network data	Reported ties	Named friends, trusted partners, peer reports
Direct relational records	Recorded interactions or relational events	Emails, social media links, direct contact records
Indirect observational records	Relational traces or proxies	Shared hometown, co-mention, co-appointment, co-participation


military connections, factional politics, political-business relations, and conflicts using indirect observational records (Shih, Adolph and Liu, 2012; Jia, Kudamatsu and Seim, 2015; Keller, 2016; Jiang, 2018; Jost and Mattingly, 2025; Bandiera, Larreguy and Mangonnet, 2026). It also gives researchers a practical way to scale network measurement beyond small hand-coded cases, surveys, and experiments.

However, treating shared attributes or observed interactions directly as “the network” creates two measurement problems. The first is *construct validity*. A proxy may measure something other than the target relationship, making the resulting network difficult to interpret. For instance, service in the same office may reflect unintentional bureaucratic rotation rather than purposeful political coalition, while a shared surname and hometown may generate spurious family ties when the surname is common within a locality. Even when a proxy is related to the concept of interest, it may capture only one aspect of the latent tie. Measuring trust through shared battlefield experience, for example, may conflate trust arising from friendship with antagonism generated by contentious coordination. Additionally, when researchers reduce rich observational sources to a single proxy edge, they discard information in archives, digital traces, images, or other data sources that could improve accuracy and reliability (Lazer and Wojcik, 2017; Brady, 2019). The second problem is *measurement error*. Proxy networks may contain both spurious ties and missing ties, and these errors are unlikely to be random. Formal, public, and well-documented relationships are easier to observe than covert, informal, private, or failed relationships. This selective visibility can also affect downstream analysis if the probability of observing a tie is related to other variables in the

analysis. Increasing sample size or adding more proxies does not by itself solve this problem.

The methodological challenge, then, is to recover theoretically important latent ties from observational data while preserving the multiplex nature of social relationships. This paper addresses that challenge by proposing **MuLTI**, a **M**ulti-indicator **L**atent **T**ie **I**nference framework for estimating multiplex latent ties from observational data. MuLTI treats social ties as latent dyadic variables and combines three types of information: observed relational traces as *noisy indicators*, shared attributes and other dyadic characteristics as *latent-structure covariates*, and variables that affect indicator visibility or documentation as *indicator-model covariates*. A loading structure, specified using substantive knowledge, maps observed indicators to latent tie types. Some indicators serve as anchors for one designated tie type, while others may load on multiple latent dimensions. For each dyad, MuLTI estimates the probability of each latent tie type, as depicted in Figure 1.

Figure 1: Multi-indicator Latent Tie Inference (MuLTI) Framework



Note: The framework separates observed indicators, tie-formation covariates, and measurement covariates. Indicators may be anchors that are theoretically specific to one latent tie type or non-anchor indicators that may be informative about multiple tie types. The target quantity is the dyad-by-tie posterior probability.

Our approach contributes to a growing methodological literature on recovering networks from noisy or indirect measurements (McCormick and Zheng, 2015; Newman, 2018; Peixoto, 2018; De Bacco et al., 2023). Existing work shows that network structure can be recovered from aggre-

gated relational data, error-prone reports, and noisy observations. MuLTI extends this approach to a setting that is common in political science, where researchers care about latent relationships and observe traces that vary in their specificity to those relationships and how likely they are to be documented. The framework makes these measurement assumptions explicit by specifying how each indicator maps to latent tie types and by separately modeling factors that affect indicator visibility. This allows researchers to combine heterogeneous relational evidence without treating any single proxy as the network or requiring direct classification of every dyad, while scaling estimation to large observational networks under a common set of assumptions and model structure.

We first develop the formal framework and establish identification of the estimand in Section 2. Section 3 explains the estimation procedure. In Section 4, simulation experiments show that MuLTI is most useful when no single observed proxy is sufficiently informative. In these settings, MuLTI substantially outperforms both single-proxy benchmarks and simple combinations of multiple proxies by pooling information across indicators, covariates, and latent network structure. Section 5 applies MuLTI to 4,424 Chinese Communist Party (CCP) elites who participated in the Chinese Communist Revolution (1927–1949), covering 9,783,676 dyads. We examine whether wartime relationships are associated with political vulnerability decades later during the Cultural Revolution (1966–1976), and whether supportive and adversarial relationships show different patterns.

2 The latent structure and indicator model

2.1 Framework

We observe dyadic data on N actors and $D = N(N - 1)/2$ undirected dyads. For each dyad d , let

$$L_d = (L_{d1}, \dots, L_{dR}) \in \{0, 1\}^R, \quad R \geq 1$$

denote a vector of latent tie types. The component $L_{dr} = 1$ indicates that dyad d carries latent tie type r . The tie types are distinct in the sense that each represents a conceptually different latent

relation motivated by the substantive context studied; they need not be mutually exclusive.

For each dyad, we observe a vector of binary indicators

$$Y_d = (Y_{d1}, \dots, Y_{dS}) \in \{0, 1\}^S, \quad S \geq R$$

where each indicator Y_{ds} records the presence or absence of an observable trace that is informative about one or more of the latent tie types.

We partition the indicators into anchor indicators and non-anchor indicators,

$$Y_d = (A_d, B_d).$$

Based on substantive knowledge, anchor indicators are assumed to be informative about only one designated tie type. Non-anchor indicators may be informative about multiple tie types. For each tie type r , let

$$A_r = (A_{r1}, \dots, A_{rJ_r})$$

denote the block of anchors assigned to tie type r , and let $B = (B_1, \dots, B_M)$ denote the non-anchor indicators.

We use two sets of dyad-level covariates. The first, $x_d^{(L)}$, enters the latent structure model and captures features relevant to tie formation (for example, attribute similarities between the two actors). The second, $x_d^{(Y)}$, enters the indicator model and captures features relevant to the measurement process (for example, the documentedness or visibility of a dyad's interactions).

Latent structure model. The latent structure model specifies the distribution of the latent tie vector L_d ,

$$\pi_d(\ell) = P(L_d = \ell \mid x_d^{(L)}, z_i, z_j), \quad \ell \in \{0, 1\}^R,$$

where z_i, z_j are actor-level latent positions, distinct from the latent tie vector L_d . We adopt a latent space specification (Hoff, Raftery and Handcock, 2002) in which $z_i^{(r)} \in \mathbb{R}^{k_r}$ are Euclidean

coordinates for actor i under tie type r in a latent space of dimension k_r .² The marginal probability of tie type r on dyad $d = (i, j)$ takes the logit-linear form

$$\text{logit } P(L_{dr} = 1 \mid x_d^{(L)}, z_i^{(r)}, z_j^{(r)}) = \gamma_{0r} + \gamma_{lr}^\top x_d^{(L)} - \|z_i^{(r)} - z_j^{(r)}\|. \quad (1)$$

The latent structure model is statistical, not causal. It specifies the distribution of latent tie configurations that, together with the indicator model, yields the conditional probabilities $\tau_d(\ell)$. The latent positions capture residual relational structure beyond what is explained by observed covariates.

Indicator model. The indicator model specifies the conditional distribution of the observed indicators given the latent tie configuration,

$$P(Y_d = y \mid L_d = \ell, x_d^{(Y)}).$$

A convenient parametric specification is

$$P(Y_{ds} = 1 \mid L_d = \ell, x_d^{(Y)}) = \sigma \left(\alpha_s + \delta_s^\top x_d^{(Y)} + \sum_{r=1}^R \mathbf{Q}_{sr} \beta_{sr} \ell_r \right), \quad (2)$$

where σ is the logistic link and $\mathbf{Q} \in \{0, 1\}^{S \times R}$ is a loading map specified using substantive knowledge. If $\mathbf{Q}_{sr} = 0$, indicator s is not allowed to depend on tie type r . For anchor indicators, the corresponding row of $\mathbf{Q}$ contains exactly one nonzero element; for non-anchor indicators, the corresponding row contains multiple nonzero elements.

Estimand. The main quantity of interest is the conditional probability of each latent tie type given the observed indicators and covariates,

$$\tau_{dr} = P(L_{dr} = 1 \mid Y_d, x_d^{(L)}, x_d^{(Y)}),$$

²Other specifications of the latent structure model are possible, and the identification results in Section 2.3 place no parametric restriction on this component.

or, more generally, the full conditional distribution over latent tie configurations,

$$\tau_d(\ell) = P(L_d = \ell \mid Y_d, x_d^{(L)}, x_d^{(Y)}).$$

Combining the latent structure model and the indicator model,

$$\tau_d(\ell) = \frac{\pi_d(\ell) P(Y_d \mid L_d = \ell, x_d^{(Y)})}{\sum_{\ell' \in \{0,1\}^R} \pi_d(\ell') P(Y_d \mid L_d = \ell', x_d^{(Y)})}.$$

Section 2.3 shows that this conditional distribution is identified under the assumptions stated in Section 2.2, together with a minimum anchor requirement.

2.2 Assumptions

We require five assumptions on the latent structure and indicator model. These assumptions support the identification results developed in Section 2.3 and ensure that the latent tie types have well-defined substantive meaning. These assumptions and results are stated pointwise in the covariate value $x = (x^{(L)}, x^{(Y)})$. For readability, we suppress this conditioning notation below. When covariates are continuous, repeated observations at any exact value of x are generally unavailable, and estimation therefore relies on the finite-dimensional restrictions imposed by the latent structure and indicator models.

Assumption 1 (Positive latent support). *At the conditioning value x ,*

$$\pi_x(\ell) = P(L = \ell \mid x) > 0 \quad \text{for every } \ell \in \{0, 1\}^R.$$

If assumption 1 fails for some ℓ , only the parameters and probabilities associated with the supported latent configurations are identified.

Assumption 2 (Exclusive anchor blocks). *For each tie type r , there exists an anchor block $A_r = (A_{r1}, \dots, A_{rJ_r})$ with $J_r \geq 3$, such that every anchor in A_r depends only on L_r , not on L_{-r} ,*

conditional on x :

$$P(A_{rj} = 1 \mid L = \ell, x) = P(A_{rj} = 1 \mid L_r = \ell_r, x).$$

Assumption 3 (Conditional independence of anchors). *Conditional on L and x , all anchors are mutually independent:*

$$P(A = a \mid L = \ell, x) = \prod_{r=1}^R \prod_{j=1}^{J_r} P(A_{rj} = a_{rj} \mid L_r = \ell_r, x).$$

The anchors must represent distinct information channels, not mechanically duplicated measurements of the same underlying interaction. Three coding variants of a single observed event do not constitute three independent anchors.

Assumption 4 (Anchor monotonicity and nondegeneracy). *For each anchor A_{rj} , define*

$$q_{rj1}(x) = P(A_{rj} = 1 \mid L_r = 1, x), \quad q_{rj0}(x) = P(A_{rj} = 1 \mid L_r = 0, x).$$

Assume

$$0 < q_{rj0}(x) < q_{rj1}(x) < 1.$$

The strict inequality in assumption 4 fixes the meaning of the state $L_r = 1$ as “tie present.” For the indicator model (2), we additionally impose $\beta_{sr} \geq 0$ for every non-anchor indicator s and every tie type r with $\mathbf{Q}_{sr} = 1$. This restriction is not required for the anchor-based identification argument, but ensures that all indicators loading on tie type r share a common direction of association with the latent state, aligning the substantive meaning of $L_r = 1$ across anchors and non-anchors.

Assumption 5 (Local independence of non-anchor indicators). *Conditional on L and x , the non-anchor indicators are independent of the anchors and of each other:*

$$P(B = b, A = a \mid L = \ell, x) = P(A = a \mid L = \ell, x) \prod_{m=1}^M P(B_m = b_m \mid L = \ell, x).$$

2.3 Anchor-Based Identification of Conditional Latent Tie Probabilities

The anchor-based identification argument proceeds in four steps. Our identification strategy requires each latent tie type to have at least three conditionally independent exclusive anchors satisfying the assumptions stated in Section 2.2. The argument adapts latent-class identification techniques (Goodman, 1974; Allman, Matias and Rhodes, 2009; Xu, 2017) to the indicator structure introduced in Section 2.1. First, Lemma 1 considers a single binary latent tie type and shows that three conditionally independent exclusive anchors are sufficient to identify the latent prevalence and the class-conditional anchor response probabilities. This result is a specialization of standard latent-class identification results for conditionally independent binary indicators. Second, Proposition 1 uses the exclusive anchor blocks to identify the joint latent tie distribution and the anchor likelihood. Third, Proposition 2 incorporates non-anchor indicators and identifies the full class-conditional indicator likelihood. Fourth, Corollary 1 combines the latent distribution and indicator likelihood to identify the target conditional latent tie probabilities.

Lemma 1 (Single-tie identification from exclusive anchors). *Fix a covariate value x . Under Assumptions 1–4 for a single tie type r , suppose the anchor block $A_r = (A_{r1}, \dots, A_{rJ_r})$ contains at least three binary anchors. Then the joint distribution of any three anchors in A_r conditional on x identifies the latent prevalence*

$$P(L_r = 1 \mid x)$$

and the anchor response probabilities

$$P(A_{rj} = 1 \mid L_r = \ell, x), \quad j = 1, 2, 3, \quad \ell \in \{0, 1\}.$$

Proof. See Appendix A.1. □

If an anchor block contains more than three anchors, any three anchors satisfying the assumptions are sufficient to identify the latent prevalence and the response probabilities for those anchors. The remaining anchors are identified from their means and covariances with the already identified

anchors. Additional anchors overidentify the single-tie measurement model and provide information for estimation and model checking. We next move from one latent tie type to the full latent tie vector. The key observation is that each anchor block yields a full-rank two-column measurement matrix, and the joint anchor measurement matrix is the Kronecker product of these matrices.

Proposition 1 (Joint identification from exclusive anchor blocks). *Fix a covariate value x . Under Assumptions 1–4, the joint latent tie distribution*

$$\pi_x(\ell) = P(L = \ell \mid x), \quad \ell \in \{0, 1\}^R,$$

is identified from the joint distribution of the anchor indicators $A \mid x$. The anchor likelihood

$$P(A = a \mid L = \ell, x)$$

is also identified for every anchor pattern a and every latent configuration ℓ . Consequently, the marginal latent tie probabilities

$$P(L_r = 1 \mid x)$$

are identified for every $r = 1, \dots, R$.

Proof. See Appendix A.2. □

Proposition 1 identifies the latent distribution and the likelihood contribution of the anchor indicators. We next incorporate the non-anchor indicators. Because non-anchor indicators may load on multiple tie types, they do not define the latent tie dimensions by themselves. Once the anchor structure identifies the latent distribution and anchor likelihood, the non-anchor indicators can be used to identify the remaining part of the indicator likelihood.

Proposition 2 (Identification of the full indicator likelihood). *Under the conditions of Proposition 1, suppose Assumption 5 also holds. Then, for every non-anchor indicator B_s and every latent*

configuration $\ell \in \{0, 1\}^R$, the class-conditional response probability

$$\theta_{s\ell}(x) = P(B_s = 1 \mid L = \ell, x)$$

is identified. Consequently, the full indicator likelihood

$$P(Y = y \mid L = \ell, x) = P(A = a, B = b \mid L = \ell, x)$$

is identified for every $y = (a, b)$ and every $\ell \in \{0, 1\}^R$.

Proof. See Appendix A.3. □

Combining the identified latent distribution with the identified full indicator likelihood gives the target conditional latent tie probabilities.

Corollary 1 (Identification of conditional latent tie probabilities). *Under the conditions of Proposition 1 and Proposition 2, the conditional distribution*

$$\tau(\ell \mid y, x) = P(L = \ell \mid Y = y, x)$$

is identified for every y and every $\ell \in \{0, 1\}^R$. Consequently, the tie-specific conditional probabilities

$$\tau_r(y, x) = P(L_r = 1 \mid Y = y, x)$$

are identified for every $r = 1, \dots, R$.

Proof. See Appendix A.4. □

Corollary 1 is the identification result for the estimand defined in Section 2.1. For dyad d , evaluating the identified conditional distribution at $Y = Y_d$ and $x = x_d$ gives

$$\tau_d(\ell) = \tau(\ell \mid Y_d, x_d), \quad \tau_{dr} = \sum_{\ell: \ell_r=1} \tau_d(\ell).$$

Under the anchor conditions above, including at least three conditionally independent exclusive anchors for each latent tie type, the results identify the latent tie distribution, the class-conditional indicator likelihood, and the conditional latent tie probabilities. They do not require identification of latent positions. Under the latent-space specification in (1), the latent positions are auxiliary parameters used to impose a parsimonious structure on $\pi_d(\ell)$ in estimation. Translations, rotations, and reflections of the latent positions leave pairwise distances unchanged and therefore do not affect τ_{dr} .

The preceding results apply to substantive tie types with exclusive anchor blocks. In some applications, non-anchor indicators may also share residual dependence not attributable to any substantively defined tie type. One modeling extension is to include an additional nuisance latent dimension, on which non-anchor indicators are permitted to load through Q , but for which no anchor block is specified. Such a nuisance dimension can absorb residual dependence among non-anchor indicators and prevent this variation from being forced onto the substantive tie types. This extension does not change the identification argument for the anchored substantive tie types in Proposition 1. However, the nuisance dimension itself is not identified by the anchor structure; its interpretation and identification depend on the loading pattern and additional restrictions imposed on the non-anchor indicators.

3 Estimation

The identification results in Section 2.3 characterize the conditional latent tie distribution and class-conditional indicator likelihoods recovered by the anchor and indicator structure at a covariate value x . For finite samples and continuous covariates, we estimate these objects through the parametric latent-structure and indicator models below.

Let

$$\Theta = \{\gamma_{0r}, \gamma_{lr}, z_i^{(r)}; \alpha_s, \beta_{sr}, \delta_s\}$$

denote the full parameter vector, indexed by tie type r , indicator s , and actor i . The latent structure model is given by (1). The indicator model is given by (2), with the pre-specified loading

restrictions encoded by $\mathbf{Q}$.

For estimation, we use the working restriction that the components of L_d are conditionally independent given the covariates and latent positions. Thus

$$\pi_d(\ell) = \prod_{r=1}^R \pi_{dr}(\ell_r), \quad \pi_{dr}(\ell_r) = P\left(L_{dr} = \ell_r \mid x_d^{(L)}, z_i^{(r)}, z_j^{(r)}\right). \quad (3)$$

This restriction is not part of the nonparametric identification argument in Section 2.3, which treats $\pi_x(\ell) = P(L = \ell \mid x)$ as an unrestricted joint distribution over latent tie configurations. Instead, (3) is a working parametric restriction that makes the latent structure model parsimonious and estimable.³ For numerical stability, we center the latent positions and constrain their average squared radius. Centering selects a representative location from the translation-equivalent class, while the radius constraint stabilizes optimization.

3.1 Observed-data likelihood

For dyad $d = (i, j)$, the complete-data likelihood factors as

$$p_{\Theta}(Y_d, L_d = \ell \mid x_d^{(L)}, x_d^{(Y)}, z_i, z_j) = \pi_d(\ell; \Theta) P_{\Theta}(Y_d \mid L_d = \ell, x_d^{(Y)}),$$

where

$$\pi_d(\ell; \Theta) = P_{\Theta}(L_d = \ell \mid x_d^{(L)}, z_i, z_j).$$

Because L_d is unobserved, estimation is based on the observed-data likelihood obtained by summing over all latent tie configurations:

$$p_{\Theta}(Y_d \mid x_d^{(L)}, x_d^{(Y)}, z_i, z_j) = \sum_{\ell \in \{0,1\}^R} \pi_d(\ell; \Theta) P_{\Theta}(Y_d \mid L_d = \ell, x_d^{(Y)}).$$

³This restriction should not be read as saying that tie types are unrelated in the observed population. Dyads may carry multiple tie types simultaneously, and fitted tie probabilities may be associated through shared covariates, correlated latent positions across tie types, or common actor-level structure. The restriction is only that residual dependence among the components of L_d is not modeled after conditioning on the covariates and latent positions included in the latent structure model.

The dyad-level observed-data log-likelihood contribution is

$$\mathcal{L}_d(\Theta) = \log \sum_{\ell \in \{0,1\}^R} \pi_d(\ell; \Theta) P_{\Theta}(Y_d \mid L_d = \ell, x_d^{(Y)}), \quad (4)$$

and the full objective is

$$\mathcal{L}(\Theta) = \sum_{d=1}^D \mathcal{L}_d(\Theta). \quad (5)$$

In implementation, we compute (4) by enumerating all $\ell \in \{0, 1\}^R$, computing the corresponding joint log probabilities,

$$a_d(\ell; \Theta) = \log \pi_d(\ell; \Theta) + \log P_{\Theta}(Y_d \mid L_d = \ell, x_d^{(Y)}),$$

and applying the log-sum-exp operation:

$$\mathcal{L}_d(\Theta) = \log \sum_{\ell \in \{0,1\}^R} \exp\{a_d(\ell; \Theta)\}.$$

We then optimize $\mathcal{L}(\Theta)$ using stochastic gradient methods. Exact enumeration is feasible because R is small. When R is large, the inner sum over $\{0, 1\}^R$ could be approximated by variational or sampling-based methods, but such approximations are not required in the settings studied here.

Given parameter values Θ , the model-implied conditional probability of latent configuration ℓ for dyad d is

$$\tau_d(\ell; \Theta) = P_{\Theta}(L_d = \ell \mid Y_d, x_d^{(L)}, x_d^{(Y)}) = \frac{\pi_d(\ell; \Theta) P_{\Theta}(Y_d \mid L_d = \ell, x_d^{(Y)})}{\sum_{\ell' \in \{0,1\}^R} \pi_d(\ell'; \Theta) P_{\Theta}(Y_d \mid L_d = \ell', x_d^{(Y)})}. \quad (6)$$

After obtaining $\hat{\Theta}$, we estimate the tie-specific conditional latent tie probability by

$$\hat{\tau}_{dr} = \sum_{\ell: \ell_r=1} \tau_d(\ell; \hat{\Theta}).$$

This optimization is related to expectation-maximization (EM) but is not implemented as an

EM algorithm. In EM, the conditional latent-configuration weights are computed at the current parameter value and held fixed while the expected complete-data log likelihood is maximized. Here, we optimize the observed-data log likelihood in (5) directly using stochastic gradient methods. The score for one dyad can be written as

$$\nabla_{\Theta} \mathcal{L}_d(\Theta) = \sum_{\ell \in \{0,1\}^R} \tau_d(\ell; \Theta) \nabla_{\Theta} \log p_{\Theta}(Y_d, L_d = \ell \mid x_d^{(L)}, x_d^{(Y)}, z_i, z_j),$$

so the gradient has a conditional-weighted complete-data form. The key distinction is that these posterior weights vary with Θ during direct optimization of the observed-data likelihood.

3.2 Stratified mini-batch optimization

For networks with hundreds or thousands of actors, the number of dyads can be large, and most dyads have no positive anchor indicators. We therefore use stratified mini-batches to increase the frequency of informative dyads during stochastic optimization of the observed-data likelihood. Rather than sampling dyads uniformly, we divide them into groups based on their observed anchor patterns and sample separately from these groups when constructing each mini-batch. This ensures that dyads with positive anchor indicators appear more frequently in the optimization than they would under uniform sampling. The strata are defined using observed anchor patterns because anchors are the indicators that carry tie-specific information by design. In each training epoch, we perform approximately D/B stochastic updates. Each update uses one mini-batch of size B .

Let $h(d) \in \{1, \dots, H\}$ denote the stratum of dyad d , and let N_h be the number of dyads in stratum h . A mini-batch draw first selects a stratum with probability p_h , then selects a dyad uniformly from that stratum. Thus the selection probability for dyad d is

$$s_d = \frac{p_{h(d)}}{N_{h(d)}}.$$

The corresponding inverse-probability weight is

$$w_d = \frac{1}{Ds_d}.$$

For a batch of size B , the inverse-probability weighted mini-batch objective is

$$\mathcal{L}^{\text{batch}}(\Theta) = \frac{1}{B} \sum_{d \in \text{batch}} w_d \log \sum_{\ell \in \{0,1\}^R} \pi_d(\ell; \Theta) P_{\Theta}(Y_d \mid L_d = \ell, x_d^{(Y)}). \quad (7)$$

Because $w_d = 1/(Ds_d)$, the expectation of (7) over the sampling design equals the average full-sample observed-data log likelihood $D^{-1}\mathcal{L}(\Theta)$. The weighted mini-batch objective therefore targets the same full-sample empirical objective.⁴

Given this weighting scheme, a key implementation choice is how to define the strata $h(d)$. We consider two broad approaches. The first stratifies dyads finely by anchor patterns. In one version, the anchor-pattern sampler uses the binary pattern over all anchors observed in the sample. This gives the greatest control over rare but informative patterns, such as dyads with multiple positive anchors. The cost is that the number of possible patterns grows exponentially in the number of anchor indicators. With $J_A = \sum_r J_r$ anchors, there are up to 2^{J_A} possible patterns. In sparse networks, only a subset of these patterns will be observed, and some observed patterns may contain very few dyads. Uniformly sampling across such fine strata can repeatedly select the same rare observations, increasing gradient variance and making updates sensitive to idiosyncratic dyads. A lower-dimensional variant stratifies dyads using one anchor indicator per latent tie type. This simplified anchor-pattern sampler preserves tie-specific anchor information but requires at most 2^R strata rather than up to 2^{J_A} strata. It is useful when the full anchor-pattern table is too sparse.

The second approach uses coarse binary strata. One version separates dyads with at least one positive anchor from dyads with no positive anchors. An even coarser version separates dyads with

⁴In implementation, we normalize the inverse-probability weights within each batch to stabilize gradient magnitudes $\tilde{w}_d = \frac{w_d}{B^{-1} \sum_{b \in \text{batch}} w_b}$. The resulting objective is a self-normalized importance-sampling version of (7). This normalization stabilizes stochastic gradients and does not change the limiting target as B grows, but it introduces the usual finite-batch bias of self-normalized importance sampling.

at least one positive indicator, anchor or non-anchor, from dyads with no positive indicators. These coarse samplers are preferable when even the simplified anchor-pattern table is sparse or when positive anchors are extremely rare. They give less control over the composition of informative dyads, but avoid repeatedly oversampling tiny fine-grained strata.

4 Simulation Experiments

We use simulation experiments to evaluate the finite-sample behavior of MuLTI under controlled data-generating conditions. The simulations focus on recovery of the latent ties L_{dr} , evaluated using the estimated posterior tie probabilities $\hat{\tau}_{dr}$.

The baseline design uses $N = 300$ actors, $R = 2$ latent tie types, two covariates in each of the latent-structure and indicator models, three anchors per tie type, and three non-anchor indicators. The baseline data-generating process induces an average latent tie prevalence of approximately 0.10 across tie types. Latent ties are generated from the latent-space structure in (1). The simulation design allows tie types to share latent geometry, so that different latent ties may be statistically associated through common actor-level structure. Conditional on covariates and latent positions, $L_{d1}, \dots, L_{dR}$ are drawn independently. Observed indicators are generated from the logistic measurement model in (2). Anchor indicators load on exactly one tie type. Non-anchor indicators may load on multiple tie types, with association strengths that may differ across tie types. For each simulated case, we fit MuLTI and estimate tie probabilities $\hat{\tau}_{dr}$ using the observed-data likelihood procedure described in Section 3, with the simplified anchor-pattern sampler unless otherwise noted.

We compare MuLTI to three anchor-based benchmarks. The single-anchor benchmark uses one anchor indicator for tie type r as the score for L_{dr} . This represents the common empirical strategy of treating one observed trace as the network measure. The anchor-count benchmark aggregates the fitted anchors for tie type r :

$$s_{dr}^{\text{count}} = \frac{1}{K_r} \sum_{k=1}^{K_r} Y_{drk}^{\text{anchor}},$$

where K_r is the number of fitted anchors for tie type r . With three fitted anchors, this score takes

values 0, 1/3, 2/3, and 1, so dyads with three positive anchors are ranked above dyads with two, one, or zero positive anchors. This benchmark uses repeated noisy measurements of the same latent tie but does not estimate covariate effects, latent geometry, or non-anchor loadings. The any-anchor benchmark collapses the same information into a binary score,

$$s_{dr}^{\text{any}} = \mathbf{1} \left\{ \sum_{k=1}^{K_r} Y_{drk}^{\text{anchor}} > 0 \right\}.$$

It ranks all dyads with at least one positive anchor above dyads with no positive anchors, but does not distinguish dyads with one, two, or three positive anchors.

The primary evaluation metrics are average precision (AP) and area under the receiver operating characteristic curve (AUC), computed separately by tie type. AP is the main metric because latent ties are sparse in many designs and the object of interest is a posterior tie probability used to rank dyads by their likelihood of carrying a tie. In this setting, AP is directly sensitive to performance among the highest-ranked dyads and is informative about recovery under class imbalance. AUC is reported as a complementary ranking metric because it measures the model’s ability to separate tied from untied dyads over all possible thresholds and is less sensitive to the marginal prevalence of the tie. For latent tie recovery, AP and AUC are computed by comparing $\hat{\tau}_{dr}$ and the benchmark scores to the true simulated L_{dr} . We first compute both metrics separately for each latent tie type and then average across tie types within each simulation case. Results report the average of these tie-level metrics across simulation runs.

4.1 Simulation Designs

Table 2 summarizes the simulation experiments reported in this section, including anchor informativeness, anchor count, network tie sparsity, and loading-map misspecification.

Table 2: Simulation Experiments

Experiment	Quantity varied	Purpose
A	Anchor strength	Sensitivity to anchor informativeness
B	Number of anchors	Sensitivity to the number of anchors used in estimation
C1	Tie prevalence	Sensitivity to network sparsity under fine stratification
C2	Tie prevalence	Sensitivity to network sparsity under coarse stratification
D	Loading map used in fitting	Robustness to loading-map misspecification

Anchor strength. Experiment A varies the informativeness of anchor indicators. For an anchor indicator $s \in \mathcal{A}_r$ assigned to tie type r , we define anchor strength on the probability scale as

$$\Delta_{sr}(x_d^{(Y)}) = \Pr(Y_{ds} = 1 \mid L_{dr} = 1, x_d^{(Y)}) - \Pr(Y_{ds} = 1 \mid L_{dr} = 0, x_d^{(Y)}).$$

For simulation summaries, we average this quantity over dyads, anchors, and tie types:

$$\bar{\Delta} = \frac{1}{R} \sum_{r=1}^R \frac{1}{|\mathcal{A}_r|D} \sum_{s \in \mathcal{A}_r} \sum_{d=1}^D \Delta_{sr}(x_d^{(Y)}).$$

This quantity is used only to describe the simulation design. In empirical applications, the latent ties are unobserved, so the corresponding true probability contrast is not directly observed. The anchor-strength experiment uses values approximately equal to


$$\bar{\Delta} \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}.$$

The baseline value used in later experiments is approximately $\bar{\Delta} = 0.50$.

Figure 2 shows that MuLTI improves as anchor indicators become more informative and consistently outperforms the anchor-based benchmarks. The gap is largest at low and moderate anchor strength, where individual anchors are noisy and the model can combine information from multiple anchors, covariates, non-anchor indicators, and latent geometry. The differences are especially pronounced for AP because latent ties are sparse in the baseline design. As anchor strength increases, the anchor-count benchmark approaches MuLTI performance. The coarser any-anchor

and single-anchor benchmarks generally underperform in AP, although their AUC gaps relative to MuLTI are smaller at high anchor strength.

Figure 2: Experiment A, Anchor Strength



Note: Points report averages after computing AP and AUC separately by tie type.


Number of anchors. For Experiment B, we generate a larger set of candidate anchors under a common data-generating process and vary the number included in estimation. This experiment evaluates how additional anchor information improves latent tie recovery.⁵

Figure 3 shows that the performance of MuLTI improves as the number of fitted anchors increases. Performance rises most between one and three anchors. Gains continue beyond five anchors but become more incremental. The anchor-count benchmark also improves as more anchors are available, approaching MuLTI at high anchor counts. The single-anchor benchmark is flat by construction. The any-anchor benchmark improves initially but then deteriorates in AP as the number of anchors grows, because the probability of at least one positive anchor increases for

⁵The identification result in Section 2.3 requires at least three anchors per tie type. Therefore, the one- and two-anchor specifications are not identified by the anchor argument. We include them as diagnostic cases because they are relevant to empirical settings with limited data availability. In these cases, the likelihood-based estimator can still produce model-based conditional probability estimates using the parametric latent-structure and measurement restrictions.

both tied and untied dyads, making the binary any-positive score less informative.

Figure 3: Experiment B, Number of Anchors



Note: Points report averages after computing AP and AUC separately by tie type.

Network sparsity. Experiment C varies network sparsity by changing the mean latent tie prevalence induced by the data-generating process. The experiment evaluates prevalence levels approximately equal to


$$\{0.01, 0.05, 0.1, 0.2, 0.3, 0.5, 0.7\}.$$

We report results under both the simplified anchor-pattern sampler and the coarse any-anchor sampler.

Figure 4 shows that recovery is hardest when ties are extremely rare. MuLTI performance improves quickly as prevalence increases from 0.01 to 0.10, with more gradual gains at higher prevalence levels. In the extremely rare case, MuLTI performs similarly to the anchor-count benchmark in AP. At higher prevalence levels, however, MuLTI consistently outperforms the anchor-based benchmarks, with the largest gains in sparse and moderately sparse settings. The AUC results are comparatively stable across prevalence levels, indicating that overall pairwise separability between tied and untied dyads changes less with sparsity. AP shows more variation because precision is

sensitive to false positives among the highest-ranked dyads. The simplified anchor-pattern sampler and the coarse any-anchor sampler produce very similar patterns, suggesting that the coarse sampler may be adequate in these simulation settings and may perform slightly better when the network is extremely sparse.

Figure 4: Experiment C, Network Sparsity



Note: Points report averages after computing AP and AUC separately by tie type.

Loading-map misspecification. Experiment D evaluates robustness to misspecification of the pre-specified loading map $\mathbf{Q}$. Data are generated under the correctly specified baseline loading map, but the model is fit using alternative loading maps that introduce different errors in the as-

sumed relationship between indicators and latent tie types. The misspecification cases, summarized in Table 3, include assigning anchors to the wrong tie type, coding anchors as non-anchors by loading on multiple tie types, coding non-anchors as anchors by loading on only one tie type, and combining anchor and non-anchor misspecification.

Figure 5 shows that MuLTI is most sensitive to misspecification that assigns anchor indicators to the wrong tie type. Performance declines modestly when one anchor is misassigned and drops more sharply when one anchor for each tie type is misassigned. In contrast, allowing anchors to load on additional tie types or restricting non-anchors to a single tie type has little effect in this design. The combined case reduces performance relative to the baseline, but less than systematic anchor misclassification alone. Even under these misspecification scenarios, MuLTI generally outperforms the anchor-based benchmarks.


Table 3: Loading-map Specifications in Experiment D

Label	Fitted loading-map specification
Baseline	Correct loading map
Wrong anchor	One anchor is assigned to the wrong tie type
Wrong anchors	One anchor for each tie type is assigned to the wrong tie type
Anchor as non-anchor	One anchor per tie type is allowed to load on multiple tie types
Non-anchor as anchor (strong signal)	Non-anchor indicators are forced to load only on the tie type with the stronger true association
Non-anchor as anchor (weak signal)	Non-anchor indicators are forced to load only on the tie type with the weaker true association
Combined	Anchor misclassification plus non-anchors restricted to their stronger-signal tie type

Note: For non-anchor indicators, strong and weak signal refer to the relative size of the true loading on each tie type in the data-generating process.

To summarize, the simulations show where MuLTI provides the largest gains over simpler anchor-based measures. The gains are largest when anchors are noisy or limited in number because the model can combine information from multiple indicators, covariates, and latent network structure that the benchmark measures ignore. When anchors are abundant and highly informative, simpler anchor-count measures approach MuLTI’s performance. Recovery is also difficult when ties are extremely rare, although MuLTI’s advantage becomes clearer as ties become sufficiently

Figure 5: Experiment D, Loading-map Misspecifications



Note: Points report averages after computing AP and AUC separately by tie type.

prevalent to provide information for estimation. Finally, the loading-map experiments highlight an important caveat: correctly assigning anchors to latent tie types matters, and assigning anchors to the wrong tie type is the most damaging misspecification considered here.

Additional simulations are reported in the appendix. Appendix B reports scaling experiments that vary network size, the number of tie types, and covariate dimensions. It also reports performance using the best-threshold F_1 score, the harmonic mean of precision and recall evaluated at the cutoff that maximizes F_1 within a given simulation case. Because this threshold is chosen ex post, we treat F_1 as a supplementary metric.

5 Empirical Application: Wartime Networks among Chinese Revolutionary Elites

We apply MuLTI to estimate latent relationships among former rebels in the Chinese Communist Revolution (1927–1949). This case illustrates the central measurement problem motivating the paper. Politically meaningful relationships among historical elites are rarely observed directly, but they often leave partial, noisy, and conceptually different traces in biographical records. A binary coding of shared hometown, shared military unit, or direct textual mentions of other actors in biographical records would miss the multifaceted relationships. MuLTI allows us to use the rich observed data to estimate probabilities of distinct latent relationships among revolutionary elites.

In this application, we study whether different types of wartime relationships are associated with individual outcomes during the Cultural Revolution (1966–1976). We distinguish between supportive and adversarial relationships because they capture two substantively different ways in which interpersonal ties may matter for political elites. Supportive relationships can provide access to resources such as mentorship, assistance, trust, and protection that individuals can draw on when needed (Lin, 2001). Adversarial relationships capture conflict, antagonism, betrayal, and punitive interactions that may constitute relational liabilities (Labianca and Brass, 2006). Estimating these dimensions separately allows us to examine whether supportive and adversarial wartime relationships have different associations with later outcomes that would be obscured by a single undifferentiated network measure. For comparison, we also use a conventional binary coworking network based on overlapping periods of organizational service, which prior research has used as a proxy for social ties among political elites (e.g., Shih, 2008; Jia, Kudamatsu and Seim, 2015; Keller, 2016; Jiang, 2018; Jost and Mattingly, 2025). We use this network as a benchmark to assess the gains from distinguishing these types of wartime relationships.

5.1 Latent Tie Dimensions, Indicators, and Covariates

In this application, we estimate two anchored substantive relational dimensions and one unanchored residual nuisance dimension. The *supportive* dimension captures relations such as trust, affinity, friendship, mutual support, mentorship, recruitment, and cooperative political or military association. The *adversarial* dimension captures disagreement, antagonism, demotion, purge, downgrading, and other punitive interactions. The *residual nuisance dimension* captures systematic dependence in observed traces that is not cleanly interpretable as either supportive or adversarial. The residual dimension is a model-dependent nuisance component used to absorb residual dependence among non-anchor indicators, so we do not interpret its posterior probabilities substantively.

To construct observed indicators of these latent relationships, we draw on the *Historical Biographical Database of the People’s Liberation Army* (HBDPLA), a new biographical archive that we compiled from official biographical sources, military reference volumes, and online encyclopedia entries.⁶ The database includes personal characteristics such as names, birth and death years, hometowns, education, work experience, and full biographical texts compiled from multiple sources. Our analysis uses 4,424 individuals in HBDPLA whose records can be matched to the *Organizational History of the CCP People’s Liberation Army* (General Political Department of the People’s Liberation Army, China, 1994).

We first partition each HBDPLA biography into wartime and postwar sections and use only the wartime portions to estimate the wartime networks. This restriction prevents post-victory appointments, later political conflicts, and other postwar information from entering the measurement of wartime relationships. We process the wartime biographies using `gpt-5.4-nano-2026-03-17`, with reasoning effort set to `xhigh`, to extract binary dyad-level indicators. These indicators capture friendly or private contact, blood kinship, marriage-established kinship, recruitment into the army or party, political purge, direct political disagreement, demotion, betrayal, policy opposition,

⁶The full dataset contains information on 7,798 individuals who joined the Chinese Communist army during the revolution and became part of the elite, defined here as individuals who held positions at the regiment level or above. Appendix C.1 provides details on the construction of the dataset.

joint combat, non-demotion appointment, joint social-movement participation, and command hierarchy. Separately, we use the *Organizational History of the CCP People's Liberation Army* (General Political Department of the People's Liberation Army, China, 1994) to construct an indicator for overlapping organizational service, defined by whether two individuals served concurrently in the same organizational unit. This indicator enters MuLTI as an observed indicator and also serves as the conventional binary network measure used for comparison. Table 4 reports the loading map and observed counts of the indicators used in the application.

The loading map imposes substantive restrictions on how each observed indicator can inform the latent relational dimensions. Some indicators are treated as anchors because, based on their substantive meaning, they are informative about only one substantive relational dimension, either supportive or adversarial. Friendly or private contact, blood kinship, marriage-established kinship, and recruitment into the army or party anchor the supportive dimension. Political purge, direct political disagreement, demotion, betrayal, and policy opposition anchor the adversarial dimension. Other indicators are allowed to load on multiple dimensions.

The latent structure model covariates include dyadic features associated with the probability that a latent relationship exists. These include shared hometown, common activist movement experience, common Whampoa Military Academy experience, shared experience in the Soviet Union (USSR), common Kuomintang experience, and common political-cadre experience. These covariates differ from the observed indicators because they are not treated as relational traces modeled by the indicator model. Unlike the indicator loadings, which are constrained so that an observed indicator provides positive evidence for the latent dimensions on which it loads, the coefficients on latent structure model covariates are not constrained to be positive. The indicator model covariates capture factors that affect whether a latent relationship is likely to appear in the observed indicators. These include military rank difference, whether one member of the dyad died during the war, birth-year difference, and the total length of the two biographies. These covariates address selective visibility in the archival recording of indicators.

Table 4: MuLTI Indicator Loading Map and Observed Indicator Counts for Wartime Networks

Observed indicator	Supportive	Adversarial	Residual	Observed indicator counts
Friendly or private contact	✓	✗	✗	521
Blood kinship	✓	✗	✗	49
Marriage-established kinship	✓	✗	✗	43
Recruitment into army or party	✓	✗	✗	791
Political purge	✗	✓	✗	370
Direct political disagreement	✗	✓	✗	379
Demotion	✗	✓	✗	120
Betrayal	✗	✓	✗	98
Policy opposition	✗	✓	✗	141
Joint combat	✓	✗	✓	1,023
Non-demotion appointment	✓	✓	✗	1,920
Joint social-movement participation	✓	✗	✓	602
Command hierarchy	✓	✓	✓	3,877
Overlapping organizational service	✓	✓	✓	25,079

Note: Bolded indicators are anchor indicators. ✓ indicates an allowed loading in the loading map; ✗ indicates a fixed-zero loading. Observed indicator counts report the number of dyads with an observed occurrence of each relational indicator. All indicators except overlapping organizational service are extracted from the wartime portions of HBDPLA biographies. Overlapping organizational service is constructed from the *Organizational History of the CCP People's Liberation Army* (General Political Department of the People's Liberation Army, China, 1994). The residual dimension is included for measurement purposes and is not interpreted substantively.

5.2 Estimated Wartime Ties and Subsequent Political Vulnerability

We estimate the model on 9,783,676 undirected dyads using the observed-data likelihood procedure described in Section 3. Estimation uses a coarse stratified mini-batch sampler that separates dyads with at least one observed indicator occurrence, whether anchor or non-anchor and regardless of latent tie dimension, from dyads with no observed indicator occurrences. For each dyad, the model returns estimated probabilities for the supportive, adversarial, and residual dimensions. Table 5 summarizes the distribution of these dyad-level estimated probabilities. The estimated probabilities are generally low for both substantively interpreted dimensions. Approximately 99% of dyads have a supportive tie probability below 0.25, and approximately 97% have an adversarial tie probability below 0.25. Only 0.60% of dyads have a supportive tie probability above 0.50, while 0.32% have an adversarial tie probability above 0.50. Dyads with very high estimated probabilities are rare: 0.14% have a supportive tie probability of at least 0.75, and 0.29% have an adversarial tie

Table 5: Distribution of Dyad-Level Estimated Tie Probabilities by Latent Dimension

Est. Prob.	Supportive		Adversarial		Residual	
	<i>N</i>	%	<i>N</i>	%	<i>N</i>	%
< 0.01	17,365	0.18	1,794	0.02	11,994	0.12
[0.01, 0.25)	9,665,745	98.79	9,444,693	96.54	9,724,779	99.40
[0.25, 0.50)	41,610	0.43	306,301	3.13	20,999	0.21
[0.50, 0.75)	44,848	0.46	2,934	0.03	243	0.00
≥ 0.75	14,108	0.14	27,954	0.29	25,661	0.26

Note: Entries report the number and percentage of dyads in each estimated-probability interval for all 9,783,676 undirected dyads among the 4,424 elites. Percentages are reported within each latent dimension. Residual denotes the unanchored nuisance dimension, which is not interpreted substantively.

probability of at least 0.75.

Next, we derive individual-level network measures by using the estimated tie probabilities as edge weights and calculating probability-weighted degree for each individual. These probability-weighted degree measures summarize each individual’s estimated exposure to the corresponding relational dimension. Using estimated tie probabilities as edge weights retains probabilistic information about dyadic tie presence without imposing an arbitrary binary classification. Let $\hat{\tau}_{ij}^C$ denote the estimated probability of a supportive tie between individuals i and j , $\hat{\tau}_{ij}^A$ the estimated probability of an adversarial tie, and $\hat{\tau}_{ij}^R$ the estimated probability of a residual association. We aggregate these dyad-level probabilities to the individual level by summing over all potential ties of each individual:

$$\text{Wartime Supportive Exposure}_i = \sum_{j \neq i} \hat{\tau}_{ij}^C,$$

$$\text{Wartime Adversarial Exposure}_i = \sum_{j \neq i} \hat{\tau}_{ij}^A,$$

$$\text{Residual Association}_i = \sum_{j \neq i} \hat{\tau}_{ij}^R.$$

For the cohort comparisons and subsequent regression analyses, we normalize each probability-weighted degree by the number of potential ties, $N - 1$. For example, normalized supportive

exposure degree centrality is

$$\widetilde{\text{Supportive}}_i = \frac{1}{N-1} \sum_{j \neq i} \hat{\tau}_{ij}^C,$$

with analogous measures for adversarial exposure and the residual dimension. The normalization places the measures on the $[0, 1]$ scale while preserving the relative ordering of individuals.

We construct the corresponding measure for the conventional coworking network in the same way. Let W_{ij} equal one if individuals i and j have an observed coworking tie and zero otherwise. Coworking degree is $\sum_{j \neq i} W_{ij}$, and normalized coworking degree centrality is


$$\widetilde{\text{Coworking}}_i = \frac{1}{N-1} \sum_{j \neq i} W_{ij}.$$

Before analyzing the association between estimated wartime relationships and mortality during the Cultural Revolution, we first describe how individual network position varies across revolutionary cohorts. Figure 6 compares the conventional coworking network with the estimated supportive and adversarial networks by year of joining the CCP and year of joining the rebel army. The three measures show different patterns. Degree centrality in the conventional coworking network is close to zero for most cohorts and shows relatively little systematic variation. Estimated supportive exposure is relatively stable across cohorts, although it is slightly higher among the earliest and latest CCP entry cohorts. Adversarial exposure shows a clearer cohort pattern: individuals who entered the CCP or rebel army earlier tend to have higher estimated adversarial exposure, with average adversarial exposure generally declining among later cohorts. Appendix C.2 provides additional descriptive analyses of the association between estimated supportive and adversarial exposure across individual characteristics.


Next, we examine whether wartime network position is associated with mortality during the Cultural Revolution. We restrict the analysis to individuals who survived until the beginning of the Cultural Revolution. We estimate the following specification:

$$Y_i = \alpha + Z_i' \beta + X_i' \gamma + \lambda_{h(i)} + \delta_{b(i)} + \eta_{e(i)} + \kappa_{p(i)} + \varepsilon_i,$$

Figure 6: Wartime Network Position across Revolutionary Cohorts



(a) Year of joining the CCP



(b) Year of joining the rebel army

Note: Each row compares three individual-level network measures across historical cohorts. The left panel reports normalized degree centrality in the conventional binary coworking network, the middle panel reports normalized probability-weighted degree in the estimated supportive network, and the right panel reports normalized probability-weighted degree in the estimated adversarial network. Bars report the mean for individuals in each cohort. Error bars report 95% confidence intervals. Cohorts are defined by year of joining the CCP and year of joining the rebel army, respectively.

where Y_i is an indicator equal to one if individual i died during the Cultural Revolution and zero otherwise. Z_i contains the wartime network measures of interest, and X_i is a vector of individual-level controls. The terms $\lambda_{h(i)}$, $\delta_{b(i)}$, $\eta_{e(i)}$, and $\kappa_{p(i)}$ denote hometown, year-of-birth, year-of-military-enlistment, and year-of-joining-the-CCP fixed effects, respectively. We estimate three specifications that differ in the network measures included in Z_i . We estimate each specification using both a linear probability model estimated by ordinary least squares (OLS) and a logistic regression. For the logistic regressions, the same linear predictor enters the logistic link function.

Table 6: Wartime Social Ties and Deaths during the Cultural Revolution

	(1)	(2)	(3)	(4)	(5)	(6)
Coworking network	4.187 (2.614)			41.816* (20.754)		
Adversarial Exposure		0.499 (0.476)	0.857+ (0.442)		16.547* (6.546)	21.079** (6.507)
Supportive Exposure			1.350 (1.191)			18.109 (13.291)
Model	OLS	OLS	OLS	Logit	Logit	Logit
Other Controls	✓	✓	✓	✓	✓	✓
Hometown FE	✓	✓	✓	✓	✓	✓
Year of Birth FE	✓	✓	✓	✓	✓	✓
Year of Military Enlistment FE	✓	✓	✓	✓	✓	✓
Year of Joining CCP FE	✓	✓	✓	✓	✓	✓
R ² / Pseudo-R ²	0.416	0.415	0.416	0.263	0.269	0.271
Observations	1843	1843	1843	991	991	991

Note: Coworking, Adversarial Exposure, and Supportive Exposure denote normalized degree centrality in the conventional coworking network, the MuLTI-estimated adversarial network, and the MuLTI-estimated supportive network, respectively. Standard errors, clustered by year of military enlistment, are reported in parentheses. Pseudo-R² values are reported for the logit models. FE = fixed effects. + $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 6 reports the estimation results. Columns (1)–(3) report OLS estimates, and Columns (4)–(6) report logit estimates. The OLS specifications use 1,843 individuals. The logit sample is smaller, with 991 observations, because fixed-effect groups with no within-group variation in the outcome are dropped from the estimation. We proceed in three steps. Columns (1) and (4) include normalized degree centrality in the conventional coworking network. Columns (2) and (5) replace this measure with estimated adversarial exposure, while Columns (3) and (6) include estimated adversarial and supportive exposure jointly. Residual association is included throughout as a control variable.

The estimated associations are positive across all three network measures, but their statistical precision differs. The conventional coworking measure is positively associated with mortality in both models, although the coefficient is statistically significant only in the logit specification. Adversarial exposure also has a positive coefficient throughout. In the OLS models, the coefficient on adversarial exposure is not statistically significant when entered alone but is marginally significant at the 10% level when supportive exposure is included. In the logit models, adversarial exposure is

statistically significant both when entered alone, at the 5% level, and when supportive exposure is included jointly, at the 1% level. Supportive exposure is also consistently positive and has a sizable point estimate, but it is estimated less precisely and is not statistically significant in either model. Taken together, these results suggest that the association with later mortality is more clearly linked to adversarial exposure than to supportive exposure. Separating wartime ties into distinct relational dimensions therefore reveals differences that are obscured by the conventional coworking measure alone.

The analysis does not provide conclusive evidence about which type of wartime relationship is associated with mortality during the Cultural Revolution, nor do we develop a theory that predicts how different wartime relationships should matter for political vulnerability during later periods of political upheaval. The results nevertheless provide some evidence that adversarial exposure accumulated decades earlier may be associated with subsequent vulnerability. The conventional coworking network also shows a positive association with mortality, but it cannot distinguish whether that association reflects supportive or adversarial relationships. MuLTI allows us to separate these dimensions and estimate their associations with the outcome independently.

6 Conclusion and Discussion

This paper proposes MuLTI as a measurement framework for settings in which theoretically important network ties are latent, multiplex, and observed only through noisy traces. By separating latent ties from their observed indicators, the framework improves on proxy-based network measurement and makes the assumptions behind network construction explicit. This approach is especially useful for historical and observational settings where the relationships of theoretical interest cannot be directly elicited but leave multiple imperfect traces in the available record.

MuLTI improves on single-anchor, any-anchor, and anchor-count benchmarks by pooling information across multiple indicators, covariates, and latent geometry, with the largest gains in average precision when individual indicators are noisy and ties are relatively sparse. The sparsity experiments show that recovery is more difficult when ties are extremely rare, but that MuLTI

generally improves on anchor-based benchmarks once there is enough signal to learn about dyads. The loading-map experiment also points to an important caution: the framework depends heavily on the correct specification of anchor indicators. Misclassifying anchors is more damaging than allowing non-anchor indicators to load imperfectly or coding anchors as non-anchors. This shows that the loading map is not a purely technical choice, but a part of the substantive measurement design.

The application to CCP revolutionary elites illustrates how MuLTI can improve on conventional binary network measures in practice. We draw on the HBDPLA and the *Organizational History of the CCP People's Liberation Army* to construct multiple dyad-level indicators of wartime relationships. MuLTI uses these indicators to estimate dyad-level probabilities of social supportive and adversarial ties. We then aggregate the estimated probabilities into individual-level measures of wartime social supportive and adversarial exposure and compare them with normalized degree centrality in a conventional coworking network. We study how these wartime network positions are associated with mortality during the Cultural Revolution among individuals who survived to the beginning of that period. The results provide suggestive evidence that wartime adversarial exposure is more clearly associated with later mortality than wartime supportive exposure, while the conventional coworking measure cannot distinguish between these relational dimensions. The application shows how MuLTI can turn heterogeneous biographical traces into distinct probabilistic network measures that can be used in downstream empirical analysis.

MuLTI provides a measurement strategy for settings where relationships of interest are not directly observed but multiple imperfect sources of information are available. Rather than asking researchers, human coders, or AI systems to classify every dyad directly, the framework separates the extraction of observable relational traces from the inference of latent relationship types. Researchers specify the latent ties of interest, how observed indicators map to those ties, and which factors affect whether relational evidence is observed; these assumptions are then applied consistently across the network. This separation is especially useful for large observational datasets, where the number of possible dyads grows quickly and direct classification becomes difficult to

scale. For example, the 4,424 individuals in our application generate more than 9.7 million dyads. With MuLTI, human or AI coding can be used to extract observable traces, while MuLTI combines those traces to estimate latent tie probabilities through a common measurement model. This approach provides a structured alternative to asking an AI system to make direct judgments about latent relationships.

The framework should be used with several cautions. First, MuLTI relies on a loading map specified using substantive knowledge, and the interpretation of each latent tie dimension depends on the validity of the anchor indicators assigned to it. Researchers should justify anchor choices, report the loading map, and conduct sensitivity analyses when alternative mappings are plausible. Second, the latent structure model captures statistical associations, not causal processes, so its coefficients should not be used to draw causal conclusions about tie formation without further assumptions. Third, the posterior tie probabilities are estimated quantities, not directly observed ties. Researchers should therefore distinguish uncertainty about latent tie presence from the binary coding used in conventional network measures.

The broader contribution is to shift political network construction from *proxy coding* toward *measurement modeling*. Rather than treating a single observed relation as the network, MuLTI uses multiple imperfect traces to estimate theoretically defined latent ties. As political scientists gain access to richer archival, textual, administrative, and digital data, this measurement perspective provides a way to use those sources without collapsing complex social relationships into a single proxy edge.

References

- Abramson, Scott F, David B Carter and Luwei Ying. 2022. “Historical border changes, state building, and contemporary trust in Europe.” *American Political Science Review* 116(3):875–895.
- Allman, Elizabeth S., Catherine Matias and John A. Rhodes. 2009. “Identifiability of Parameters in Latent Structure Models with Many Observed Variables.” *The Annals of Statistics* 37(6A).
- Bandiera, Antonella, Horacio Larreguy and Jorge Mangonnet. 2026. “Family ties, social control, and authoritarian distribution to elites.” *American Political Science Review* 120(1):226–244.
- Bastian, Mathieu, Sebastien Heymann and Mathieu Jacomy. 2009. “Gephi: an open source software for exploring and manipulating networks.” *Proceedings of the International AAAI Conference on Web and Social Media* 3(1):361–362.
- Brady, Henry E. 2019. “The challenge of big data and data science.” *Annual Review of Political Science* 22(1):297–323.
- Brierley, Sarah and Noah L Nathan. 2021. “The connections of party brokers: Which brokers do parties select?” *The Journal of Politics* 83(3):884–901.
- Campbell, David E. 2013. “Social networks and political participation.” *Annual Review of Political Science* 16(1):33–48.
- Cranmer, Skyler J, Philip Leifeld, Scott D McClurg and Meredith Rolfe. 2017. “Navigating the range of statistical tools for inferential network analysis.” *American Journal of Political Science* 61(1):237–251.
- Cruz, Cesi, Julien Labonne and Pablo Querubin. 2020. “Social network structures and the politics of public goods provision: evidence from the Philippines.” *American Political Science Review* 114(2):486–501.

De Bacco, Caterina, Martina Contisciani, Jonathan Cardoso-Silva, Hadiseh Safdari, Gabriela Lima Borges, Diego Baptista, Tracy Sweet, Jean-Gabriel Young, Jeremy Koster, Cody T Ross, Richard McElreath, Daniel Redhead and Eleanor A Power. 2023. “Latent network models to account for noisy, multiply reported social network data.” *Journal of the Royal Statistical Society Series A: Statistics in Society* 186(3):355–375.

URL: <https://academic.oup.com/jrsssa/article/186/3/355/7030939>

General Political Department of the People’s Liberation Army, China. 1994. *Zhongguo Gongchandang Zhongguo Renmin Jiefang Jun Zuzhi Shi Ziliao (Organizational History of the Communist Party of China - Chinese People’s Liberation Army)*. Changzheng Publishing House.

Goodman, Leo A. 1974. “Exploratory Latent Structure Analysis Using Both Identifiable and Unidentifiable Models.” *Biometrika* 61(2):215–231.

Hall, Peter A and Michèle Lamont. 2013. “Why social relations matter for politics and successful societies.” *Annual Review of Political Science* 16(1):49–71.

Hicken, Allen. 2011. “Clientelism.” *Annual review of political science* 14(1):289–310.

Hoff, Peter D, Adrian E Raftery and Mark S Handcock. 2002. “Latent Space Approaches to Social Network Analysis.” *Journal of the American Statistical Association* 97(460):1090–1098.

URL: <http://www.tandfonline.com/doi/abs/10.1198/016214502388618906>

Jia, Ruixue, Masayuki Kudamatsu and David Seim. 2015. “Political selection in China: The complementary roles of connections and performance.” *Journal of the European Economic Association* 13(4):631–668.

Jiang, Junyan. 2018. “Making bureaucracy work: Patronage networks, performance incentives, and economic development in China.” *American Journal of Political Science* 62(4):982–999.

Jost, Tyler and Daniel Mattingly. 2025. “Networks of coercion: Military ties and civilian leadership challenges in China.” *American Journal of Political Science* . eprint:

<https://onlinelibrary.wiley.com/doi/pdf/10.1111/ajps.12971>.

URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/ajps.12971>

Keller, Franziska Barbara. 2016. “Moving beyond factions: using social network analysis to uncover patronage networks among Chinese elites.” *Journal of East Asian Studies* 16(1):17–41.

Labianca, Giuseppe (Joe) and Daniel J. Brass. 2006. “Exploring the Social Ledger: Negative Relationships and Negative Asymmetry in Social Networks in Organizations.” *Academy of Management Review* 31(3):596–614.

Larson, Jennifer M. 2021. “Networks of conflict and cooperation.” *Annual Review of Political Science* 24(1):89–107.

Larson, Jennifer M and Janet I Lewis. 2020. “Measuring networks in the field.” *Political Science Research and Methods* 8(1):123–135.

Lazer, David. 2011. “Networks in political science: Back to the future.” *PS: Political Science & Politics* 44(1):61–68.

Lazer, David and Stefan Wojcik. 2017. Political Networks and Computational Social Science. In *The Oxford Handbook of Political Networks*, ed. Jennifer Nicoll Victor, Alexander H. Montgomery and Mark Lubell. Oxford University Press pp. 115–130.

Lin, Nan. 2001. *Social Capital: A Theory of Social Structure and Action*. Structural Analysis in the Social Sciences Cambridge: Cambridge University Press.

Masterson, Daniel. 2024. “Refugee networks, cooperation, and resource access.” *American Political Science Review* 118(3):1398–1414.

McCormick, Tyler H. and Tian Zheng. 2015. “Latent Surface Models for Networks Using Aggregated Relational Data.” *Journal of the American Statistical Association* 110(512):1684–1695.

Publisher: ASA Website eprint: <https://doi.org/10.1080/01621459.2014.991395>.

URL: <https://doi.org/10.1080/01621459.2014.991395>

- McPherson, Miller, Lynn Smith-Lovin and James M Cook. 2001. “Birds of a feather: Homophily in social networks.” *Annual review of sociology* 27(1):415–444.
- Newman, M. E. J. 2018. “Network structure from rich but noisy data.” *Nature Physics* 14(6):542–545.
URL: <https://www.nature.com/articles/s41567-018-0076-1>
- Organization Department of the Chinese Communist Party. 2000. *Zhongguo Gongchandang Zuzhishi Ziliao (Organizational History of CCP)*. Zhonggong Dangshi Chubanshe (Chinese Communist Party History Press).
- Peixoto, Tiago P. 2018. “Reconstructing Networks with Unknown and Heterogeneous Errors.” *Physical Review X* 8(4):041011.
URL: <https://link.aps.org/doi/10.1103/PhysRevX.8.041011>
- Perliger, Arie and Ami Pedahzur. 2011. “Social network analysis in the study of terrorism and political violence.” *PS: Political Science & Politics* 44(1):45–50.
- Ravanilla, Nico, Renard Sexton and Dotan Haim. 2022. “Deadly populism: How local political outsiders drive Duterte’s war on drugs in the Philippines.” *The Journal of Politics* 84(2):1035–1056.
- Shami, Mahvish. 2012. “Collective action, clientelism, and connectivity.” *American Political Science Review* 106(3):588–606.
- Shih, Victor C. 2008. *Factions and Finance in China: Elite Conflict and Inflation*. New York: Cambridge University Press.
- Shih, Victor, Christopher Adolph and Mingxing Liu. 2012. “Getting ahead in the communist party: explaining the advancement of central committee members in China.” *American Political Science Review* 106(1):166–187.

- Ward, Michael D, Katherine Stovel and Audrey Sacks. 2011. “Network analysis and political science.” *Annual Review of Political Science* 14:245–264.
- Wood, Elisabeth Jean. 2008. “The social processes of civil war: The wartime transformation of social networks.” *Annual Review of Political Science* 11(1):539–561.
- Xu, Gongjun. 2017. “Identifiability of Restricted Latent Class Models with Binary Responses.” *The Annals of Statistics* 45(2).
- Zaytsev, Dmitry G, Valentina V Kuskova, Gregory S Khvatsky and Anna A Sokol. 2023. “Searching for coherence in a fragmented field: Temporal and keywords network analysis in political science.” *Network Science* 11(1):113–142.

Online Appendix for

MuLTI: A Framework for Measuring Latent Ties from Rich but Noisy Observational Data

Table of Contents

A	Proofs	A-1
A.1	Proof of Lemma 1	A-1
A.2	Proof of Proposition 1	A-2
A.3	Proof of Proposition 2	A-3
A.4	Proof of Corollary 1	A-4
B	Additional Latent Tie Prediction Simulation and Evaluation	A-5
B.1	Best-Threshold F_1 Evaluation	A-6
C	Additional Details on the Application to Chinese Revolutionary Elites	A-8
C.1	The Dataset – HBDPLA	A-8
C.2	Additional Descriptive Patterns in Estimated Wartime Networks	A-9

A Proofs

A.1 Proof of Lemma 1

Proof. With J binary anchors, the unrestricted two-class model has $1 + 2J$ free parameters: one latent prevalence $p = P(L = 1)$, J false-positive probabilities $q_{j0} = P(A_j = 1 \mid L = 0)$, and J true-positive probabilities $q_{j1} = P(A_j = 1 \mid L = 1)$. The observed joint distribution of J binary anchors has $2^J - 1$ free cell probabilities. A necessary condition for point identification is therefore

$$2^J - 1 \geq 1 + 2J,$$

which fails for $J = 1$ and $J = 2$, and first holds at $J = 3$.

At $J = 3$, define

$$\Delta_j = q_{j1} - q_{j0}, \quad c_{ju} = \text{Cov}(A_j, A_u),$$

and

$$\mu_3 = E \left[\prod_{h=1}^3 (A_h - E[A_h]) \right].$$

By conditional independence and the definition of $q_{j\ell}$,

$$E[A_j \mid L] = q_{j0} + \Delta_j L, \quad E[A_j] = q_{j0} + p \Delta_j.$$

Conditional independence of the anchors given L implies

$$c_{ju} = \text{Cov}(E[A_j \mid L], E[A_u \mid L]) = p(1 - p) \Delta_j \Delta_u.$$

Let $v = p(1 - p)$, so that $c_{ju} = v \Delta_j \Delta_u$.

For the third centered moment, write

$$A_j - E[A_j] = (E[A_j \mid L] - E[A_j]) + \varepsilon_j, \quad E[\varepsilon_j \mid L] = 0.$$

Conditional independence implies that every term in the expansion of $\prod_h (A_h - E[A_h])$ containing at least one ε_h has conditional expectation zero given L . Hence

$$\mu_3 = E \left[\prod_{h=1}^3 (E[A_h \mid L] - E[A_h]) \right] = \Delta_1 \Delta_2 \Delta_3 E[(L - p)^3].$$

For a Bernoulli random variable with success probability p , $E[(L - p)^3] = p(1 - p)(1 - 2p)$, so

$$\mu_3 = v(1 - 2p) \Delta_1 \Delta_2 \Delta_3.$$

Now form the observable ratio

$$R = \frac{\mu_3^2}{c_{12} c_{13} c_{23}} = \frac{v^2(1 - 2p)^2 \Delta_1^2 \Delta_2^2 \Delta_3^2}{v^3 \Delta_1^2 \Delta_2^2 \Delta_3^2} = \frac{(1 - 2p)^2}{v}.$$

Since $(1 - 2p)^2 = 1 - 4p(1 - p) = 1 - 4v$,

$$R = \frac{1 - 4v}{v} = \frac{1}{v} - 4, \quad \text{so} \quad v = \frac{1}{R + 4}.$$

This identifies $v = p(1 - p)$. The equation $v = p(1 - p)$ has roots

$$p = \frac{1 - \sqrt{1 - 4v}}{2} \quad \text{and} \quad p = \frac{1 + \sqrt{1 - 4v}}{2}.$$

If $\mu_3 = 0$, then $p = 1/2$. If $\mu_3 \neq 0$, the sign of μ_3 identifies the sign of $1 - 2p$, because $v > 0$ and $\Delta_1 \Delta_2 \Delta_3 > 0$ by Assumption 4. Hence the appropriate root is selected, and p is identified.

Next, using the covariance identities,

$$\Delta_1^2 = \frac{c_{12}c_{13}}{v c_{23}}, \quad \Delta_2^2 = \frac{c_{12}c_{23}}{v c_{13}}, \quad \Delta_3^2 = \frac{c_{13}c_{23}}{v c_{12}}.$$

Monotonicity gives $\Delta_j > 0$, so the positive roots identify $\Delta_1, \Delta_2, \Delta_3$. Finally,

$$q_{j0} = E[A_j] - p \Delta_j, \quad q_{j1} = q_{j0} + \Delta_j, \quad j = 1, 2, 3.$$

Thus p and the six anchor response probabilities are identified. □

A.2 Proof of Proposition 1

Proof. For each tie type r , let $\mathcal{A}_r$ denote the finite set of possible response patterns for the anchor block A_r . Define the block measurement matrix $M_r(x)$ by

$$[M_r(x)]_{a_r, \ell_r} = P(A_r = a_r \mid L_r = \ell_r, x), \quad a_r \in \mathcal{A}_r, \quad \ell_r \in \{0, 1\}.$$

By Lemma 1, the two class-conditional anchor-block distributions

$$P(A_r = a_r \mid L_r = 0, x) \quad \text{and} \quad P(A_r = a_r \mid L_r = 1, x)$$

are identified. They are distinct by monotonicity. Since each column of $M_r(x)$ is a probability vector, two distinct columns cannot be scalar multiples of one another. Therefore,

$$\text{rank}(M_r(x)) = 2.$$

Thus each block likelihood $P(A_r = a_r \mid L_r = \ell_r, x)$ is identified.

Let

$$A = (A_1, \dots, A_R)$$

denote the full vector of anchor blocks. Let g_x be the vector of observed joint anchor-pattern probabilities,

$$g_x(a) = P(A = a \mid x),$$

and let π_x be the 2^R -vector with entries

$$\pi_x(\ell) = P(L = \ell \mid x), \quad \ell \in \{0, 1\}^R.$$

By exclusivity of anchor blocks and conditional independence of anchors given L and x ,

$$P(A = a \mid L = \ell, x) = \prod_{r=1}^R P(A_r = a_r \mid L_r = \ell_r, x).$$

Since each factor on the right-hand side is identified, the anchor likelihood $P(A = a \mid L = \ell, x)$ is identified for every a and ℓ .

Stacking the joint anchor-pattern probabilities over all a gives

$$g_x = M(x)\pi_x,$$

where

$$M(x) = \bigotimes_{r=1}^R M_r(x).$$

The rank of a Kronecker product is the product of the ranks. Therefore,

$$\text{rank}(M(x)) = \prod_{r=1}^R \text{rank}(M_r(x)) = 2^R.$$

Thus $M(x)$ has full column rank and a left inverse. Hence

$$\pi_x = M(x)^+ g_x$$

for any left inverse $M(x)^+$. Therefore the full joint latent tie distribution π_x is identified.

The marginal tie probabilities follow by summing over latent configurations:

$$P(L_r = 1 \mid x) = \sum_{\ell \in \{0,1\}^R: \ell_r=1} \pi_x(\ell).$$

□

A.3 Proof of Proposition 2

Proof. By Proposition 1, the joint latent distribution π_x , the anchor likelihood $P(A = a \mid L = \ell, x)$, and the anchor measurement matrix $M(x)$ are identified. Therefore the anchor-only conditional distribution

$$W_x(a, \ell) = P(L = \ell \mid A = a, x)$$

is identified by Bayes' rule:

$$W_x(a, \ell) = \frac{P(A = a \mid L = \ell, x)\pi_x(\ell)}{P(A = a \mid x)}.$$

Equivalently, in matrix form,

$$W_x = D_g(x)^{-1} M(x) D_\pi(x),$$

where $D_g(x)$ is the diagonal matrix of observed anchor-pattern probabilities and $D_\pi(x)$ is the diagonal matrix of latent class probabilities. By positive latent support and nondegenerate anchor probabilities, the diagonal entries of both matrices are positive. Since $M(x)$ has full column rank, W_x also has full column rank.

For a fixed non-anchor indicator B_s , define the 2^R -dimensional vector

$$\theta_s(x) = (\theta_{s\ell}(x) : \ell \in \mathcal{L}), \quad \theta_{s\ell}(x) = P(B_s = 1 \mid L = \ell, x), \quad \mathcal{L} = \{0, 1\}^R.$$

For each anchor pattern a , define the observed conditional probability

$$m_s(a, x) = P(B_s = 1 \mid A = a, x).$$

By Assumption 5,

$$P(B_s = 1 \mid L = \ell, A = a, x) = P(B_s = 1 \mid L = \ell, x).$$

Therefore,

$$\begin{aligned} m_s(a, x) &= \sum_{\ell \in \mathcal{L}} P(B_s = 1 \mid L = \ell, x) P(L = \ell \mid A = a, x) \\ &= \sum_{\ell \in \mathcal{L}} W_x(a, \ell) \theta_{s\ell}(x). \end{aligned}$$

Stacking over all anchor patterns gives

$$m_s(x) = W_x \theta_s(x).$$

Since W_x has full column rank, $\theta_s(x)$ is uniquely identified.

By Assumption 5, the non-anchor indicators are mutually independent conditional on L and x . Hence

$$P(B = b \mid L = \ell, x) = \prod_{s=1}^M \theta_{s\ell}(x)^{b_s} \{1 - \theta_{s\ell}(x)\}^{1-b_s}.$$

Thus the non-anchor likelihood $P(B = b \mid L = \ell, x)$ is identified.

Assumption 5 also gives conditional independence between anchors and non-anchors given L and x . Therefore,

$$P(Y = y \mid L = \ell, x) = P(A = a, B = b \mid L = \ell, x) = P(A = a \mid L = \ell, x) P(B = b \mid L = \ell, x).$$

The first factor is identified by Proposition 1, and the second factor is identified above. Hence the full indicator likelihood $P(Y = y \mid L = \ell, x)$ is identified. $\square$

A.4 Proof of Corollary 1

Proof. By Proposition 1, $\pi_x(\ell) = P(L = \ell \mid x)$ is identified. By Proposition 2, $P(Y = y \mid L = \ell, x)$ is identified. Therefore Bayes' rule identifies

$$P(L = \ell \mid Y = y, x) = \frac{\pi_x(\ell) P(Y = y \mid L = \ell, x)}{\sum_{\ell' \in \{0,1\}^R} \pi_x(\ell') P(Y = y \mid L = \ell', x)}.$$

The tie-specific conditional probability follows by summing over latent configurations:

$$P(L_r = 1 \mid Y = y, x) = \sum_{\ell: \ell_r=1} P(L = \ell \mid Y = y, x).$$

$\square$

B Additional Latent Tie Prediction Simulation and Evaluation


This appendix reports additional latent tie prediction simulations and supplementary evaluation metrics. The experiment numbering in this appendix continues from the main article.

Network size. To assess scaling with the number of actors and dyads, we vary

$$N \in \{100, 300, 500, 700, 900\}.$$

This design changes the number of dyads while leaving the remaining baseline design fixed.

Figure B1: Experiment E, Network Size



Note: Points report averages after computing AP and AUC separately by tie type.

Number of tie types. We vary the number of latent tie types over

$$R \in \{1, 2, 3, 4\}.$$


For $R \geq 2$, the design retains the three non-anchor indicators from the baseline specification. For the single-tie case, we omit non-anchor indicators, since cross-loading is not defined when there is only one latent tie type.

Covariate dimension. Finally, we vary the dimensions of the tie-layer and measurement-layer covariates:

$$(p_L, p_Y) \in \{(0, 0), (2, 2), (6, 6), (10, 10)\}.$$


This experiment evaluates sensitivity to the number of observed dyad-level covariates included in the latent-structure and indicator models.

Figure B2: Experiment F, Number of Tie Types



Note: Points report averages across tie types and simulation replications.

Figure B3: Experiment G, Covariate Dimension




Note: Points report averages across tie types and simulation replications.

B.1 Best-Threshold F_1 Evaluation

We use best-threshold F_1 as a supplementary diagnostic. For each simulation case, we choose the cutoff that maximizes F_1 , where F_1 is the harmonic mean of precision and recall. Cutoff selection applies to MuLTI and the anchor-count benchmark, whose scores are not binary. The single-anchor and any-anchor benchmarks are binary by construction, so their F_1 scores are evaluated directly. Because the cutoff is selected using the realized outcomes, this metric should be interpreted only as an ex post summary of classification performance. Results are reported in Figure B4. MuLTI outperforms the anchor-based benchmarks in almost all cases, except for the one-anchor, extremely rare-tie, and multiple misattributed anchor cases.

Figure B4: Best-Threshold F_1 Performance Across Latent Tie Prediction Experiments



Note: Points report averages across tie types and simulation replications. F_1 is evaluated at the threshold that maximizes F_1 within each simulation case.

C Additional Details on the Application to Chinese Revolutionary Elites

C.1 The Dataset – HBDPLA

Construction and Scope Drawing on various official biographical sources, we compiled the *Historical Biographical Database of the People's Liberation Army* (HBDPLA), which encompasses approximately 7,800 Red Army officers. To ensure accuracy, we gathered up to three biographical records for each officer, cross-verifying the details provided.

The primary sources include Baidu Encyclopedia, Wikipedia, and multiple volumes of *Biographical Portrait of Figures in Party History*: (1) Profiles of the Chinese Red Army (*Zhongguo hongjun renwu zhi*), (2) Profiles of the First Front Army of the Chinese Workers' and Peasants' Red Army (*Zhongguo gongnong hongjun diyi fangmianjun renwu zhi*), (3) Profiles of the Fourth Front Army of the Chinese Workers' and Peasants' Red Army (*Zhongguo gongnong hongjun disi fangmianjun renwu zhi*), (4) Directory of Martyrs of the Fourth Front Army of the Chinese Workers' and Peasants' Red Army (*Zhongguo gongnong hongjun disi fangmianjun lieshi minglu*), (5) Biography of the Martyrs of the Third Army Corps of the Chinese Workers' and Peasants' Red Army (*Zhongguo gongnong hongjun disan juntuan ying lie zhuan*), and (6) Dictionary of the Central Revolutionary Base Area (*Zhongyang geming genju di cidian*).

The HBDPLA tracks a wide range of individual attributes — including names, birth and death years, hometown, education, cause of death, family background, social movement experience, and career trajectory. Additionally, these rich biographical texts enable the discovery of interpersonal ties and interactions among militants.




Figure C1: Biographical Archive

Directory of Martyrs of the Fourth Front Army of the Chinese Workers' and Peasants' Red Army

Sample of Ex-combatants The HBDPLA specifically documents historical figures who served at or above the *regimental level*⁷ in the CCP's military organization from its inception.⁸ This focus was facilitated by the CCP's extensive archival records and abundant military publications. We restricted the sample

⁷The military organization hierarchy was: Front (*fang'mian jun*), Field Army (*jun'tuan*), Corps (*jun*), Division (*shi*), Regiment (*tuan*), Company (*lian*), and Platoon (*pai*).

⁸Since 1927, the armed organizations of the Chinese Communist Party have had different names: the National Revolutionary Army led by the Chinese Communist Party, the Chinese Workers' and Peasants' Red Army during the first Chinese Civil War, the New Fourth Army and Eighth Route Army during the War of Resistance Against Japan,

to individuals active before the founding of the People’s Republic of China (P.R.C.), thus ensuring that each person had previously been a rebel participating in the Chinese Civil War – consistent with the basic definition of an “armed rebel.”

To eliminate duplicates, we cross-referenced name variations, aliases, and pseudonyms. As a result, each individual in the dataset corresponds to at least one official biography, verified through multiple sources. Nonetheless, even the editors of these official biographical compilations note two caveats: (1) Reliance on personal recollections: important historical details (e.g., birth/death dates, hometowns) may be inaccurate. (2) Uneven documentation: the level of detail varies significantly across individuals. Therefore, to maximize historical accuracy and completeness, we supplemented each official biography with additional archival sources and online encyclopedia entries, giving each individual up to three different biographical records.

C.2 Additional Descriptive Patterns in Estimated Wartime Networks

Figure C2 shows the association between estimated supportive and adversarial exposure across individual characteristics. For many characteristics, the positive category is relatively rare, so the zero category closely resembles the overall distribution. Supportive and adversarial exposure are positively associated for several groups, but the strength of this association varies across characteristics. The dispersion around the fitted relationship also differs across groups. Figure C7 examines the associations between the estimated supportive and adversarial exposure measures, biography length, and conventional coworking degree. Because the MuLTI measurement model explicitly adjusts for differences in documentation, the estimated supportive and adversarial exposure measures show little systematic association with either overall or wartime biography length. Both measures, however, have a modest positive association with conventional coworking degree. This pattern is expected because individuals with more coworking ties have more opportunities to develop both supportive and adversarial relationships.

and, after the outbreak of the second Chinese Civil War, the People’s Liberation Army.




Figure C2: Association between Estimated Supportive and Adversarial Exposure by Individual Characteristics

Note: The first panel shows the ungrouped scatterplot. The remaining panels plot estimated adversarial exposure against adversarial exposure separately by a binary individual characteristic. Points are binned averages, and lines show fitted linear relationships within each group.




Figure C3: Associations between MuLTI Estimates, Biography Length, and Coworking Degree

Note: The panels plot estimated supportive exposure (left column) and adversarial exposure (right column) against overall biography length (top row), wartime biography length (middle row), and conventional coworking degree (bottom row). Biography length captures the extensiveness of biographical documentation. Lines show local polynomial smooths, with shaded areas indicating 95% confidence intervals.




Figure C4: Network graph based on conventional co-work ties

Note: The graph is plotted with Gephi (Bastian, Heymann and Jacomy, 2009).




Figure C5: Three horizontally aligned plots with explanatory notes.

Note: The graph is plotted with Gephi (Bastian, Heymann and Jacomy, 2009).




Figure C6: Three horizontally aligned plots with explanatory notes.

Note: The graph is plotted with Gephi (Bastian, Heymann and Jacomy, 2009).




Figure C7: Network graph based on conventional co-work ties

Note: The graph is plotted with Gephi (Bastian, Heymann and Jacomy, 2009).