From inference to influence: applying causal game theory to complex security environments

Vonk, M.C.

Chapter 7

Conclusions and Future Work

This thesis advances the application of (strategic) causal frameworks within the context of complex security environments. It does so by surveying and developing methodological approaches that are subsequently applied to substantive policy challenges. This final chapter provides a summary of all previous chapters and thereby discusses how this thesis contributes to addressing the research questions. The conclusion then turns to address the broader implications and contributions of this research, structured around the challenges of causal inference, its policy relevance, remaining research gaps, and the thesis's theoretical, methodological, empirical, and societal significance. Finally, future research avenues are proposed.

7.1 Summary

Chapter 1: The introduction establishes the context for integrating causal inference with strategic interaction in complex security environments. It identifies a critical research gap: while game theory is widely used in strategic studies, causal inference remains notably underutilized despite its importance for evidence-based policy-making. This gap persists because complex security environments—featuring interference effects, strategic interdependence, and adaptive adversaries—violate fundamental assumptions of standard causal models. Additionally, the mathematical complexity and computational demands of existing causal methods create implementation barriers for practitioners. The chapter outlines how this dissertation addresses these challenges through an integrated framework that enables the application of causal inference in complex security environments. The research questions are structured to examine the

necessary conceptual frameworks (RQ1), develop computationally efficient methods (RQ2), and demonstrate practical applications (RQ3) in such security environments.

Chapter 2: The preliminary chapter presents the fundamentals of probability theory, graphical models, and game theory. These elements establish the foundational terminology and conceptual framework that underpin the development of the subsequent chapters.

Chapter 3: This chapter organizes fragmented causal inference methods into Pearl's causal hierarchy. The hierarchy contains three levels: association (seeing), intervention (doing), and counterfactual (imagining). Each level requires specific mathematical tools. The associational level uses Bayesian networks to model observational relationships. The interventional level employs causal Bayesian networks and do-calculus to estimate causal effects. The counterfactual level needs structural causal models for retrospective analysis. By identifying which causal concepts belong to each level, this chapter answers RQ1.1: *What fundamental causal concepts are necessary for structuring and differentiating causal relationships, particularly in the context of Pearl's causal hierarchy?* Additionally, the chapter reveals critical assumptions for each method, answering RQ1.2: *What key assumptions underpin causal inference applications across Pearl's causal hierarchy?* It shows how relaxing standard assumptions necessitates alternative causal structures that explicitly model the resulting complexities. It also distinguishes parametric methods from non-parametric methods. For contexts where the no-interference assumption fails, the chapter adopts the spatially explicit structural equation model as an alternative framework that explicitly models spillover effects. The chapter equips practitioners with decision criteria to select methods based on their policy questions.

Chapter 4: This chapter integrates causal reasoning with strategic decision-making to model complex security environments. The chapter introduces game-theoretic foundations (normal-form, extensive-form, and Bayesian games) before extending them to incorporate causal structures. Multi-agent influence diagrams combine game theory with causal graphs, enabling simultaneous modeling of strategic interactions and causal mechanisms. These form the basis of causal games, where equilibrium strategies induce causal Bayesian networks. By showing how causal games enable policy-makers to compute both strategic equilibria and causal intervention effects within a unified framework, the chapter answers RQ1.3: *What methods exist for integrating causal rea-*

soning with strategic decision-making in complex security environments, and how can they be applied? The chapter provides practical guidance on model selection based on problem structure and details the elicitation requirements for implementation. Examples from deterrence scenarios demonstrate how this framework captures the interplay between causal effects and strategic adaptation in complex security environments.

Chapter 5: The most technical contributions of this thesis are presented in this chapter, where a method is presented to approximate optimal causal interventions in hybrid Bayesian networks. The chapter first introduces an approach for approximate inference based on discretization and decision diagrams. This approach achieves over 10x speedup compared to traditional methods, while Pareto fronts reveal how practitioners can balance computational cost against inference accuracy. These visualizations show precisely how many discretization bins to use for desired accuracy levels, providing implementation guidance. This systematic analysis of the accuracy-efficiency trade-off answers RQ2.1: *How can inference be performed efficiently to accurately estimate the effects of causal interventions while maintaining computational feasibility?* Finally, this method is then embedded in a broader framework that, in combination with optimization algorithms, can approximate optimal interventions in such hybrid Bayesian networks. Empirical evidence shows that while local methods (BFGS, Powell, OnePlusOne) perform poorly, differential evolution and NGOpt consistently outperform random search, demonstrating that these problems contain sufficient structure for heuristic optimizers to exploit. This performance comparison across multiple optimization techniques answers RQ2.2: *How can optimization techniques be integrated with causal inference to optimize over causal interventions efficiently under budget constraints?*

Chapter 6: This chapter applies the previously introduced concepts and validates the scientific contributions in the context of complex security environments, thereby addressing RQ3: *How can the proposed (strategic) causal concepts be applied to complex security environments?* Two applications are examined. First, causal frameworks from Chapter 3 are applied to environmental conflict in Iraq. A causal discovery algorithm uncovers the empirical structure linking environmental factors to conflict from 294 municipal observations. The analysis reveals that soil moisture and latent energy affect conflict through demographic and agricultural mediators. Spatially explicit structural equation models estimate these effects while accounting for interference between municipalities. This application proves that causal methods can identify

actionable intervention points in conflict prevention. The second application examines hybrid threat deterrence using the causal game-theoretic framework from Chapter 4. A causal influence diagram models how states select counter-hybrid measures against cyber attacks on critical infrastructure. Integer linear programming identifies market restrictions and intelligence sharing as optimal counter-hybrid measures across 1000 scenarios. The multi-agent extension computes subgame perfect equilibria, showing that the deterrer prioritizes cost-effective deterrence over damage mitigation when adversaries act strategically. Sensitivity analysis confirms that the effectiveness of counter-hybrid measures depends primarily on adversary responsiveness. This application demonstrates how causal game theory informs security policy under strategic contestation.

7.2 Conclusions

Causal inference has evolved significantly across disciplines since David Hume's 18th-century philosophical inquiries into the nature of causation. Medicine advanced the field through the development of randomized controlled trials in the mid-20th century, establishing rigorous methods for determining treatment effects. Computer science has recently transformed causal inference from observational data, with Judea Pearl developing graphical models and the do-calculus for causal reasoning [178], Peter Spirtes pioneering constraint-based algorithms for causal discovery [231], and Elias Bareinboim advancing methods for causal inference across heterogeneous domains [28].

While these sophisticated causal methods have remained largely confined to computer science and clinical domains, strategic studies would benefit immensely from their systematic application to evaluate policy interventions. Policy-makers must distinguish between actions that genuinely reduce conflict and those that merely correlate with peaceful periods. Without causal analysis, security interventions risk squandering resources on ineffective measures or inadvertently escalating conflicts through misguided policies. Accurate prediction of causal intervention outcomes becomes particularly essential in complex security environments. In these environments, interventions trigger cascading effects that spread across regions and organizations, making it crucial for policymakers to understand both immediate outcomes and secondary consequences. Moreover, adversaries actively exploit gaps in causal understanding by adapting their tactics to circumvent interventions, shifting operations to areas where policy effects are weakest.

However, applying causal inference to complex security environments faces signifi-

cant obstacles. Standard causal models rely on assumptions that these environments often violate, especially the no-interference assumption, which requires that the intervention of one unit does not affect the outcome of another. In practice, policy interventions can cause armed groups to move to neighbouring areas, creating spillover effects that undermine these assumptions. Additionally, traditional causal frameworks struggle to model multiple strategic actors who anticipate and adapt to interventions. Adversaries observe security policies and change their tactics, altering the very causal relationships that analysts aim to understand. These technical limitations, along with the mathematical complexity of advanced causal methods, create a considerable implementation gap, preventing security practitioners from using tools that could improve their decision-making.

This dissertation addresses these challenges by extending causal inference beyond its traditional computational foundations, adapting its core concepts to be more applicable within complex security studies. By clarifying foundational concepts in both causality and game theory, it demonstrates how these analytical tools can be integrated to tackle the unique problems of complex security environments. The thesis illustrates how causal reasoning, when combined with models of strategic interaction, provides a robust framework for understanding and guiding policy interventions in environments defined by uncertainty, interdependence, and competing objectives.

The findings of this research reveal that concepts from causal inference, when properly adapted to account for the characteristics of complex security environments such as interference and strategic interactions, provide a powerful framework for addressing these environments. Specifically, the research demonstrates that while sophisticated causal methods already exist in other domains, such as spatially explicit structural equation models in ecology, they require systematic extraction, restructuring, and adaptation to become operational in strategic studies. By organizing these disparate methods within Pearl's causal hierarchy and connecting them to specific policy questions, this thesis establishes how causal tools can be systematically applied to improve decision-making in security environments characterized by uncertainty, interdependence, and competing interests.

The thesis further concludes that the implementation gap between theoretical causal methods and practical security applications can be bridged through computational innovation. The development of efficient approximation methods for optimal causal interventions in hybrid Bayesian networks demonstrates such innovation. By achieving over 10x speedup compared to traditional methods while maintaining accuracy, and by providing clear visualization of accuracy-efficiency trade-offs, these tools

enable policy-makers to make informed choices about computational resources versus analytical precision. This computational contribution transforms abstract causal theory into actionable decision support.

Finally, this thesis concludes that causal game-theoretic frameworks generate actionable policy insights when applied to real-world security contexts. The case studies on environmental conflict in Iraq and hybrid threat deterrence demonstrate that these models can effectively handle the defining features of complex security environments: spatial interference between units, strategic adaptation by multiple interdependent actors, and pervasive uncertainty. The empirical findings reveal that latent energy and soil moisture indirectly cause conflict activity through demographic and agricultural mediators, while the game-theoretic analysis identifies which characteristics of counter-hybrid measures effectively deter adversaries from conducting hybrid operations. These results demonstrate that this integrated approach enhances security policy effectiveness by anticipating strategic responses and enabling targeted interventions based on causal understanding.

In conclusion, this research provides a structured, accessible approach to the application of causal inference within strategic studies. The contributions of this thesis span multiple dimensions, each addressing distinct aspects of the challenge of applying causal inference to complex security environments. The following sections detail how this work advances the field through theoretical, methodological, empirical, and societal contributions.

Theoretical Contributions: The dissertation presents a structured framework for integrating causal inference and strategic interaction in complex security environments. Pearl's causal hierarchy, which organizes causal reasoning across three levels (association for observational relationships, intervention for effects of manipulations, and counterfactuals for retrospective analysis of alternative scenarios), provides the foundational structure. The first theoretical contribution of this dissertation systematically maps existing causal concepts and their underlying assumptions onto this hierarchy, providing practitioners with clear guidance on which causal tools are appropriate for specific security policy questions. This mapping enables more effective alignment between policy questions and analytical methods by making explicit the assumptions required at each level.

Recent work has begun integrating game-theoretic elements into causal models [91], but these approaches have remained largely bereft of real-world application due to their technical complexity. The second theoretical contribution enhances the practical accessibility of these existing methods by explicitly specifying the foundational game-

theoretic elements, clarifying their intersection with causal concepts, and detailing the model inputs and data requirements needed for implementation in security contexts. This operationalization bridges the gap between theoretical frameworks and practical application.

Methodological Contributions: To address implementation barriers, the thesis develops computational innovations that make sophisticated analytical tools practically viable for security practitioners. The primary methodological contribution is a novel approach for approximating causal interventions in hybrid Bayesian networks through discretization and knowledge compilation. This method directly addresses the computational constraints that can prevent practitioners from applying analytical tools. It demonstrates through empirical evaluation how the approach balances inference accuracy with computational feasibility. The trade-off between computational cost and accuracy is systematically analyzed and visualized through Pareto fronts, which provide practitioners with guidance on parameter selection based on their specific accuracy and resource requirements. The dissertation embeds the approximation method within an optimization framework that enables evaluation of multiple causal interventions. This framework benchmarks various optimization algorithms and allows practitioners to identify optimal interventions while respecting real-world resource limitations.

Empirical Contributions: This dissertation demonstrates practical applicability through two applications that exemplify complex security challenges. The first application analyzes environmental conflict in Iraq, applying causal discovery methods to uncover mechanisms linking environmental factors to violence. The analysis uses literature-informed variable selection to guide causal discovery and accounts for spatial interdependencies in computing causal estimates of conflict incidence across geographical units.

The second application focuses on hybrid threat deterrence, demonstrating how strategic causal models can inform counter-hybrid strategies. This case models adversarial hybrid operations using causal influence diagrams and uses integer linear programming to identify optimal counter-hybrid measures. The application extends to multi-agent settings where subgame perfect equilibria are computed to determine optimal strategies under strategic contestation, with sensitivity analysis providing insights into how different variables impact equilibrium outcomes.

Societal Contributions: These empirical applications demonstrate the thesis's broader societal and policy contribution by validating how strategic causal models inform real-world decision-making. The Iraq case study reveals how environmental

variables drive conflict through specific causal pathways that enable targeted policy responses that account for spatial spillovers between regions. The hybrid threat analysis quantifies the effectiveness of different deterrence measures, helping policymakers optimize resource allocation across defensive measures. These findings illustrate the practical value of integrating causal reasoning with strategic elements to support evidence-based policy in contested environments. Beyond security contexts, this approach applies to any domain where policymakers must navigate both causal complexity and strategic interdependence among competing actors.

While this thesis advances the integration of causal inference and strategic interaction, significant challenges remain that limit the full realization of these methods in complex security environments. The contributions presented here represent important steps forward, yet they also illuminate critical gaps that must be addressed before causal game-theoretic frameworks can reach their full potential in such security environments.

Theoretically, Pearl's causal hierarchy rests upon structural causal models where recursiveness is often assumed. However, most complex security environments exhibit inherent feedback loops and cyclical dynamics. Escalatory dynamics, where actions and counteractions spiral through cycles of increasing intensity, are fundamentally cyclic in nature and violate standard recursiveness assumptions. The theoretical foundations for integrating cyclic causal models with strategic interaction remain underdeveloped, limiting the framework's ability to capture these dynamics in strategic settings.

Methodologically, the no-interference assumption, which requires that interventions for one unit do not affect others' outcomes, is highly unlikely to hold in complex security environments where spillover effects are ubiquitous. While this thesis introduces spatially explicit structural equation models to partially address this limitation, the vast majority of causal research continues to rely on no-interference assumptions. Methods that can simultaneously handle various forms of interference, while accounting for strategic interaction, remain computationally intractable or theoretically incomplete.

Empirically, real-world security settings often suffer from limited, sensitive, and inconsistently structured data due to secrecy, classification, and covert activities. At the same time, there is a lack of rigorous empirical validation, as it is rarely possible to evaluate the effectiveness of real-world policy interventions. Most causal game-theoretic methods require complete data or rely on strong assumptions when data is missing, making them ill-suited to these challenges. Strategic concealment further

complicates both data availability and validation, reinforcing the gap between theory and practice.

7.3 Future Work

Numerous avenues for advancing the theory, methodology, and empirical validation have been explored throughout this thesis. The present discussion is limited to future directions aimed specifically at addressing the broader research gaps outlined above.

Addressing the theoretical gaps requires extending causal frameworks beyond their current acyclicity constraints. Research is now emerging that accounts for the existence of feedback loops in structural causal models [36], but more research is necessary to fully integrate these cyclic structures with strategic contestation. Future work should develop mathematical foundations that can represent escalatory dynamics and feedback loops while maintaining the analytical tractability needed for policy applications. This includes establishing identification criteria for causal effects in cyclic strategic systems and developing equilibrium concepts that account for how causal relationships evolve through repeated strategic interactions.

To overcome methodological limitations, future research must develop scalable approaches for relaxing the no-interference assumption in strategic settings. Computational tractability is particularly critical here, as approximating causal interventions under the no-interference assumption is already computationally demanding, and introducing interference effects multiplies these challenges significantly. This highlights the need for targeted research into which specific forms of interference are most relevant in complex security environments, followed by the development of specialized interference models that maintain computational tractability as a primary design constraint. Promising directions include leveraging machine learning advances to approximate complex interference patterns efficiently and creating diagnostic tools that help practitioners systematically identify which types of interference are most critical in their specific contexts.

With respect to empirical data constraints, tailored data collection strategies must be developed to accommodate the specific characteristics of complex security environments. For example, in environmental security, international efforts to compile climate-conflict indicators have enabled the environmental conflict analysis in this thesis. Similarly, advancing causal models in domains such as hybrid threat deterrence will require investment in datasets capturing cyber activities, influence operations, and multi-actor strategic behaviour. For domains where enhanced data availability

remains unlikely, greater emphasis should be placed on structured expert elicitation, including formal extraction of utilities, belief distributions, and strategic type assessments from domain experts. While methods exist for discrete settings, further work is needed to adapt these techniques for continuous domains and complex interdependencies. Regarding validation constraints, monitoring mechanisms must be built to assess the effectiveness of interventions in real-world security contexts. This should enable systematic evaluation of whether theoretical predictions translate into practical outcomes and allow for further refinement of causal game-theoretic frameworks.

Finally, a valuable direction for future research lies in the development of plug-and-play tools that enable users to apply causal game-theoretic analysis without requiring proficiency in programming or advanced formal modeling. Such tools should offer predefined templates, guided workflows, and user-friendly interfaces that allow practitioners to input domain-specific knowledge and data, while automating the underlying computational procedures. Embedding these methods within the curricula of strategic and security studies, particularly as part of quantitative methods training, would further support their diffusion. Doing so would equip future analysts and decision-makers with accessible tools for causal reasoning in strategic contexts, without placing undue technical demands on users.