Phenomenological agent-based modeling: A case study of the Dutch inland shipping sector

Master thesis submitted to Delft University of Technology

in partial fulfilment of the requirements for the degree of

MASTER OF SCIENCE

in **Engineering and Policy Analysis**

Faculty of Technology, Policy and Management

by

Eunice Koid

Student number: 5278023

To be defended in public on August 15, 2022

Graduation committee

Chairperson	: Prof.dr. M.E. Warnier,	Systems Engineering
First Supervisor	: Dr. H.G. van der Voort,	Organization & Governance
External Supervisor	: Dr. S.I. Wassenburg,	Inspectie Leefomgeving en Transport
External Supervisor	: Ir. Jasper van Vliet,	Inspectie Leefomgeving en Transport
External Supervisor	: P.P.A.B. Merkx, MSc,	Inspectie Leefomgeving en Transport

Executive Summary

Law enforcement occurs in a complex environment that contains a variety of actors that interact with one another. These interactions create emerging collective behavior over time. For example, inspectors will try to influence non-compliant actors to become compliant, while inspectees may comply or thwart inspections. Inspection agencies such as the Inspectie Leefomgeving en Transport (ILT) evaluate inspectees' adherence to regulation and aim to boost compliance within an industry. However, they have limited resources and a wide range of potential societal challenges to address. They face an *action dilemma*, having to decide on what actions to take without full knowledge of whether their actions lead to higher compliance or improved social outcomes. Previous studies of the inspection environment rely on behavioral theories to investigate the underlying motivations of inspectees' behavior. However, these theories presume inspectees' motivations and characterize them homogeneously, often assuming they have perfect rationality. This leads to an inaccurate depiction of inspectees, reducing them to one-dimensional actors when in reality, their behavior is motivated by multiple factors and can be idiosyncratic.

Data science techniques provide the opportunity to understand behavioral phenomena with data, leveraging datasets to identify statistical patterns in behavior to help inspectorates make decisions within a degree of certainty. A particular modeling technique that focuses on representing empirical data without assuming behavioral motivations is phenomenological modeling. Coupled with agent-based modeling (ABM), phenomenological modeling allows the researcher to simulate possible outcomes of observed behavior before the underlying motivations are understood. This provides insight into the macro-level behavior produced by micro-level interactions within the complex inspection environment.

This research employs a phenomenological approach to ABM to determine which inspection and enforcement strategies are effective for increasing compliance within the Dutch inland shipping sector. Therefore, the main research question addressed in this research is: *How does a phenomenological approach to agent-based modeling avoid inaccurate presumptions on inspectees' behavioral motivations and show which inspection and enforcement strategies are most effective at increasing compliance among inspectees that exhibit idiosyncratic behaviors?* Analyses of longitudinal inspection data and inspectors' qualitative data informed the development of the Phenomenological Agent-based Model for Inspections (PABMI). The PABMI models

three behavioral phenomena: peer pressure, reaction to inspection, and reaction to enforcement. Responsive enforcement and five inspection strategies were included: all random, risk-based (highest non-compliant record), risk-based (highest offense severity), mix of random and risk-based (highest non-compliant record), and mix of random and risk-based (highest offense severity). These inspection strategies and responsive regulation were simulated in four scenarios that reflect possible inspectee environments: Individualistic, Non-responsive (SIM-1); Networked, Non-responsive (SIM-2); Individualistic, Responsive (SIM-3); and Networked, Responsive (SIM-4).

The results of the PABMI show that inspections alone are not sufficient for improving compliance, contrary to preconceived notions of the ILT. To increase compliance in the sector, the ILT should seek to influence positive peer pressure and deliver enforcement interventions that reduce non-compliant behavior in a short time span. Inspectors should transition to a pure risk-based inspection strategy for non-responsive environments, targeting inspection candidates that have high non-compliant records. In a responsive environment, inspectors should use a mix of random and risk-based inspections where half of inspection candidates are identified randomly and the other half are identified based on their offense severity. In addition, inspectors should commit to a responsive enforcement strategy where the severity of the enforcement is commensurate with the severity of the offense. In addition, the ILT should devote resources to the exploration and data collection of the effects of peer pressure and responsiveness in the sector.

This research contributes to existing literature by applying two modeling techniques hardly used in the study of the inspection environment: phenomenological modeling and agent-based modeling. The method of combining statistical analysis and behavioral theories resulted in the proposed Operational Framework for a Phenomenological Approach to ABM. This framework offers an iterative operational cycle that streamlines inspectorates' model development process despite their resource constraints and bounded rationality. This offers a process to refine existing behavioral theories, data collection methods, and models so that they can be a powerful tool for assessing the effectiveness of an inspectorate's strategies. Meanwhile, this process spurs discussion within inspectorates and provides a structure for discussing unforeseen drivers of behavior. Most importantly, this research provides a societal contribution by directing the ILT towards how they can better fulfill their mission of safeguarding the sustainability of society and the environment.

Contents

Acronyms

ABM Agent-based model(ing). 2–4, 15–24, 30, 38, 42, 45, 61, 89, 90, 93–101, 103, 104

AR All random inspection strategy. 35, 44, 45, 49, 69, 72, 76, 87, 90, 186–240, 242–244, 246–253

GBTM Group-based Trajectory Modeling. 25, 26, 29, 97, 98, 101, 102, 104

ILT Inspectie Leefomgeving en Transport (Human Environment and Transport Inspectorate). ii, 1, 9, 10, 20–22, 24, 25, 28, 30, 32, 33, 38, 39, 42, 44, 61, 76, 92–104

LHS Landelijke Handhavingsstrategie (National Enforcement Strategy). 9, 10, 20, 32, 33, 36, 37, 39, 42, 43, 92, 100, 102

MRRbNC Mixed random and risk-based inspection strategy based on non-compliant record. v, vi, 35, 54–58, 60, 76, 82, 88, 128, 129, 133–142, 147–151

MRRbOS Mixed random and risk-based inspection strategy based on offense severity. v, vi, 35, 54, 58–61, 78, 80, 85, 87, 88, 91, 92, 103, 152–156, 162–166, 171–175

PABMI Phenomenological Agent-based Model for Inspections. v, 2, 3, 30, 31, 34, 35, 37–45, 49, 62–64, 76–78, 84, 90–99, 101–104

RbNC Risk-based inspection strategy based on non-compliant record. 35, 78, 80–82, 85, 88, 91, 92, 103

RbOS Risk-based inspection strategy based on offense severity. 35, 81, 87, 91

RE Responsive enforcement. 36, 37, 39, 43, 45–47, 50, 51, 53–58, 60–62, 64–69, 72–75, 77, 115–117, 124–128, 131, 132, 143–151, 155, 156, 167–175, 189, 190, 195, 196, 207–216, 220, 221, 232–237, 240–244, 248–253

SE Standard enforcement. 36, 37, 39, 46, 47, 49–58, 60–62, 64–67, 69–74, 76, 77, 112–114, 118–123, 130, 131, 133–142, 154, 155, 157–166, 193, 194, 197–206, 219, 220, 222–231, 238–240, 245–247

Nomenclature

$\%_{init,comp}$	Initial percentage of compliant inspectees
$\%_{rand}$	Percent of random inspections
μ	Logit normal mean
σ	Logit normal standard deviation
r_{peers}	Radius of peers

Chapter 1

Introduction and Literature Review

1.1 Introduction

Law enforcement occurs in a complex environment that contains a variety of actors that exist in a network, each having different self-interests and incentives. At a minimum, this network of enforcement contains inspectees, inspectors, and third parties like citizens and the media (De Bruijn et al., 2007). These groups of actors interact with one another, which gives rise to evolving collective behavior over time. For example, inspectors will always try to induce compliant behavior in inspectees, and inspectees can respond to this in a way that frustrates or complies with the desires of inspectors (De Bruijn et al., 2007).

Inspectorates like the Inspectie Leefomgeving en Transport (ILT) – the Human Environment and Transport Inspectorate in the Netherlands – are responsible for 1) evaluating inspectees' adherence to rules and regulations, 2) boosting compliance throughout the industry, and 2) safeguarding the safety and sustainability of society, environment, infrastructure, transport, and housing (Inspectie Leefomgeving en Transport, 2022). However, inspectorates have limited resources to conduct inspections, and there are an endless number of societal challenges that can be addressed (Black and Baldwin, 2010). They face what De Bruijn et al. (2007) call the *action dilemma*; inspectors must decide on what action to take, like choosing who to inspect and what intervention strategies to use, without knowing for certain whether their decision leads to higher compliance and better social outcomes. Their finite capacity and lack of access to perfect information mean that they have bounded rationality, where they are forced to make the best possible choice with the resources and information they have (Simon, 1957). Yet, choosing wrongly could have socially damaging consequences. For example, if inspectors do not conduct adequate enforcement, they risk allowing non-compliant actors to continue operating, which may lead to social and environmental damage. At the other extreme, if inspectors conduct too much enforcement, they risk wasting resources and unnecessarily limiting the operations of an industry.

Previous studies into the enforcement network have hypothesized various underlying motivations of inspectees' actions in an effort to predict behavior and devise intervention strategies. Notably, various authors such as Becker (1968), Kagan and Scholz (1980), and Ayres and Braithwaite (1992) developed typologies of compliance – theoretical frameworks that categorize actors based on their motivational drivers. While useful as a starting point to model behavior, their validity is limited. The main criticism of these typologies is that they do not fully explain observed behavior; it is impossible to model the multitude of causes that could potentially drive behavior. Therefore, models that rely on these typologies may not accurately describe the observed reality that some inspectees react to inspections in counter-intuitive and unexpected ways.

In the last two decades, the growing field of data science expands opportunities to "analyze and understand actual phenomena with data" (Hayashi, 1998, p. 41). Data science techniques have the potential to alleviate inspectors' action dilemma by leveraging data to identify historical patterns, allowing inspectorates to make informed predictions about the future within a degree of certainty (Adhikari and DeNero, 2017). Rather than hypothesizing about the causality of behavior, data science focuses on drawing conclusions based on the statistical analysis of empirical observations. A particular modeling technique that only focuses on representing empirical data is phenomenological modeling. It does not require an understanding of underlying motivations to simulate behavioral changes over time. Coupled with agent-based modeling (ABM), a phenomenological approach can provide insight into the interactions that shape a complex network of enforcement and help inspectorates find effective intervention strategies that boost compliance rates.

This research employs a phenomenological approach to ABM to help inspectorates better understand the network of enforcement and make more informed, data-driven decisions. In addition, it evaluates the usefulness of a phenomenological approach to ABM in filling knowledge gaps within the inspection environment due to insufficient data and incomplete behavioral theories. In doing so, this research assesses the trade-offs of using a phenomenological approach compared to a traditional approach to ABM to help inspectorates address their action dilemma and find effective intervention strategies that boost compliance rates.

1.1.1 Thesis Layout

Chapter 1 contains an introduction to the research and literature review. Chapter 2 discusses the methods used for the research. Chapter 3 outlines the conceptualization of the Phenomenological Agent-based Model for Inspections (PABMI) developed in this research. Chapter 4 discusses the implementation of the PABMI.

Chapter 5 and 6 presents the results of the PABMI and a discussion of the results, respectively. Finally, Chapter 7 ends with a conclusion.

1.1.2 Research Aim and Main Question

This research assesses the usefulness of a phenomenological approach to ABM for determining intervention strategies that increase compliance rates, given a variety of behaviors observed by inspectors and discovered in hard data. To this end, the main research question is:

How does a phenomenological approach to agent-based modeling avoid inaccurate presumptions on inspectees' behavioral motivations & show which inspection and enforcement strategies are most effective at increasing compliance among inspectees that exhibit idiosyncratic behaviors?

This research question contains two crucial parts. The first part emphasizes that this research evaluates the value of applying a phenomenological approach to ABM, specifically to avoid an overdependence on imprecise behavioral theories. While theories provide a good starting point to investigate a topic further, relying on them too heavily poses risks. The main risk of an overreliance on existing behavioral theories is that it presumes the underlying motivations of behavior, often in a reductive way. Therefore, models based on a reductive characterization of actors' behaviors may produce results that do not accurately describe the real-life environment in which the actors exist.

The second part of the research question highlights that this research assesses how a phenomenological approach to ABM provides insight into effective inspection and enforcement strategies given a population of inspectees that exhibit a variety of behaviors. The phenomenological approach bypasses the reductive nature of behavioral theories by focusing on statistical behavioral patterns found in quantitative and qualitative data. This integrates several behavioral observations into the model so that they can be studied in combination with each other. Finding effective strategies given a population of inspectees that act counter-intuitively and in contrast with one another helps inspectorates address their action dilemma; they can discover where to best direct their resources to achieve the best compliance outcomes with more certainty.

1.1.2.1 Research Sub-questions

To answer the main research question, the following sub-questions are addressed:

1. What are the theoretical foundations for the application of phenomenological and agent-based modeling techniques for studying compliance behavior?
2. What theories underpin the conceptualization of a phenomenological ABM of the inspection environment?
3. What data and analyses are needed to identify observed behavioral phenomena in the inspectee population?
4. How can behavioral phenomena be conceptualized into an ABM?
5. How can the effectiveness of inspection and enforcement strategies be investigated under different scenarios of the inspectee population?

1.2 Literature Review

The literature review examines existing research to understand compliance theories that already exist, identify its limitations, and inform the theoretical basis of this research. In addition, it reviews methodological approaches that have been used to study compliance behavior, including its advantages and challenges. The aim of this literature review is to determine the theoretical foundation of the use of phenomenological ABMs for studying compliance behavior. Doing so illuminates its usefulness and potential for further research upon which this thesis is built.

1.2.1 Typology of Compliance

Compliance is defined as the "acquiescence to expectations that can take a range of forms: rules, standards, proposals, entreaties, orders, [and] suggestions" (Étienne, 2010). Most notably, Kagan and Scholz (1980) present three basic typologies of compliance: 1) The amoral calculator, 2) The political citizen, and 3) The incompetent. This tripartite categorization of inspectees provides a distinction between key archetypes of inspectees' motivations, the reason behind why inspectees choose to comply or violate regulations. Said differently, these typologies are an attempt to describe patterns of causes behind a particular behavior.

The essence of these three typologies are cited and discussed extensively in other research. To illustrate, the typology of the amoral calculator is examined by Becker (1968) in his study of a utilitarian, economically-driven theory on inspectees' motivation. He proposes that for an economically-motivated inspectee, illegal action is chosen over other alternatives if the economic utility of committing a crime is greater than its costs. Amoral calculators are economical, self-interested actors, driven by what May (2005) and Mitchell (2007) refer to as the *logic of consequence*. With this logic, inspectees make decisions based on explicit calculations on how the consequences of the decision will affect their interests. This typology has been widely used in later studies to model inspectees' behavior (see examples in Stigler, 1970; Becker and Stigler, 1974; and May, 2005).

The political citizen is an inspectee who exhibits what Mitchell (2007) calls the *logic of appropriateness*. An inspectee driven by this logic is motivated by socialized and internalized norms. They often ask themselves, "what is the *right* thing to do in this situation for someone like me?" (Mitchell, 2007). These actors are primarily driven by a sense of duty and adherence to norms. Tyler (1990) indicates that political citizens may be driven either by an acceptance of an authority as legitimate or by the desire to act according to individual ethical standards. These actors adhere to laws simply because they accept them as social norms that are morally good.

The incompetent typology of compliance refers to inspectees whose behavior is caused by a lack of capacity or competence to recognize that they are not abiding by laws. Van Snellenberg and Van de Peppel (2002) argue that in some cases, organizational ineptitude can cause violations of the law. Examples of organizational ineptitude can include poor management and leadership, insufficient procedures, and lack of coordination (Van Snellenberg and Van de Peppel, 2002; De Bruijn et al., 2007). Moreover, divergent goals at different levels of an organization can lead to confusion and misalignment, making the environment too complex to determine whether compliance is achieved (Van Snellenberg and Van de Peppel, 2002; De Bruijn et al., 2007).

In addition to Kagan and Scholz's three typologies of compliance, the Tafel van Elf (Table of Eleven) is another theoretical framework that delineates typologies of inspectees' behaviors. Created by Dick Ruimschotel, it delineates the motivations behind why individuals break rules based on previous behavioral science studies (Ruimschotel et al., 1996). The framework has 11 dimensions that are categorized into two groups: 1) Spontaneous, intrinsic compliance and 2) Enforcement-driven compliance (see Table 1.1).

MOTIVATION	DESCRIPTION	PROPOSED INTERVENTION
Spontaneous, Intrinsic Behavior		
1. Knowledge of the rules	An unfamiliarity of regulations and/or a lack of clarity of regulations leads to inadvertent non-compliance.	Target communications and training about regulations.
2. Costs and benefits	Financial and intangible costs and benefits influence non-compliance. This includes the financial costs and benefits of compliance versus non-compliance and the intangible costs and benefits to public reputation.	Subsize compliant actors, price regulations, certification of good conduct.
3. Degree of acceptance	Individual acceptance of the reasonableness of the regulation.	Encourage self-regulation by placing responsibility of the policy's success on the inspectees themselves.
4. Adherence to standards	The degree of willingness of the inspectee to conform to the authority of the government.	Education on the policy's purpose.
5. Social control	The inspectees' perceived probability of positive or negative sanctions and/or backlash for their behavior by other actors in their social network, such as competitors, other companies, other non-governmental monitoring organizations, and customers.	Provide information to the social network on quality standards and how they can better identify non-compliance.
Enforcement-driven Behavior		
6. Perceived likelihood of being reported to the government	Perceived probability that a violation will be detection by parties other than the government who then report it to the government.	Set up click lines and improve accessibility of inspectorates, stimulate willingness to report.
7. Perceived likelihood of administrative or physical control	Perceived probability of inspection by the government.	Publicize information about inspection probabilities, increase inspection capacity.
8. Perceived likelihood of detection	Perceived probability of detection of violation by the government.	Publicize information about the likelihood of discovery, increase resources to increase enforcement capacity.
9. Perceived likelihood of being inspected in the future	Perceived probability of increasing inspections in the future in the event of a violation today.	Investigate those with increased risk of violation, namely, those who have violated in the past.
10. Perceived likelihood of being sanctioned	Perceived probability of a sanction if a violation is found after inspection.	Increase capacity of the Public Prosecution Service and public administration.
11. Severity of sanctions	The amount and type of sanction linked to the violation and its associated disadvantages.	Increase sanctions with more violations.

Table 1.1: Tafel van Elf
(adapted from Ministerie van Justitie, 2006).

These behavioral theories postulate how inspectees make decisions given a myriad of options, oftentimes assuming that they have *perfect rationality*. The perfect rationality model is the standard assumption of a traditional, economically-driven characterization of human behavior (Beinhocker, 2006). Perfect rationality simplifies an agent's decision-making process and behavior by assuming that they hold core values, have perfect information about their environment, and possess the capacity to make detailed calculations on which decision would help them achieve their desired outcome (Beinhocker, 2006; Marin et al., 2020). While perfect rationality is useful for simplifying human behavior and developing tractable equations upon which models can be built, it makes unrealistic assumptions about behavior. First, it assumes that inspectees prioritize one goal, like economic utility, over everything else. This depicts them as one-dimensional actors; in reality, inspectees are usually driven by more than one motivation and thus do not fall exclusively under one typology (Étienne, 2010). Oftentimes, the most accurate explanation of the drivers of action is a combination of various motivations (Alm et al., 1995; Scholz, 1997; May, 2005). Second, perfect rationality assumes that inspectees can make complex and detailed calculations, considering large amounts of information to reason which action will produce their desired outcomes according to their core beliefs. While this allows researchers to build predictive models on inspectees' behavior, it ignores the complexity of the real world; inspectees often do not have access to large swaths of data nor the ability to process it (Beinhocker, 2006). Finally, the perfect rationality model pigeonholes inspectees into acting a certain way, disregarding their agency to change over time. Because behavioral theories and typologies do not explain all the causal factors of observed behavior (Étienne, 2010), forcing inspectees into typologies limits the breadth of real-world behavioral phenomena that can be studied. If inspectees behave in a perfectly rational way, then all their actions will be deterministic (Beinhocker, 2006); every inspectee with the same amount of information in the same environment would make the same decision. This ignores the fact that inspectees could behave in unexpected and idiosyncratic ways. For example, inspectees might make mistakes in their calculations which causes them to unintentionally make a utility-minimizing decision. One inspectee might react positively to law enforcement, while another might react negatively to it. In addition, inspectees may unpredictably prioritize different goals at different points in time. Models that are based on imprecise behavioral theories run the risk of modeling behavior in unrealistic ways while ignoring key aspects of how inspectees actually behave.

1.2.1.1 Proposed Enforcement Strategies for Typologies of Compliance

Enforcement is the coercion of inspected parties to act in compliance with laws without which there is no obligation for the parties to act compliantly (Lodge and Wegrich, 2012). Enforcement strategies employed by inspectors aim to change non-compliant behavior to compliant behavior (Lodge and Wegrich, 2012). Existing literature attempts to suggest prescriptive enforcement strategies to deal with each typology of compliance. For amoral calculators, authors recommend an economic-related strategy; either deterrence (increasing the likelihood of getting caught or increasing the cost of violation) or the provision financial incentives for compliance is appropriate. The goal of this strategy is to change behavior through economic means, as it assumes that decreasing utility in response to non-compliant behavior will disincentivize bad behavior (Becker, 1968; Burby and Paterson, 1993; Gray and Scholz, 1993; Van Snellenberg and Van de Peppel, 2002). However, subsequent empirical studies disprove the argument that heavier fines lead to more compliance (see Grasmick and Green, 1980; Paternoster et al., 1982; Tsebelis, 1990; and Sherman, 1993). Moreover, research by Tsebelis (1991) and Tsebelis (1993) shows that raising fines in some situations can have no impact or even a negative impact on compliance.

In other cases, Langbein and Kerwin (1985) discover that decreasing fines can potentially lead to more compliance. These counter-intuitive and divergent results led to further exploration of other qualitative enforcement strategies by other scholars. For political citizens and incompetent actors, Kagan and Scholz (1980) suggest a cooperative and educational strategy, respectively. A cooperative strategy maintains a balanced relationship of negotiation and collaboration between inspectees and inspectors. An educational strategy aims to build the capacity of inspected parties to understand regulations and cultivate the processes and procedures necessary for compliance (Kagan and Scholz, 1980).

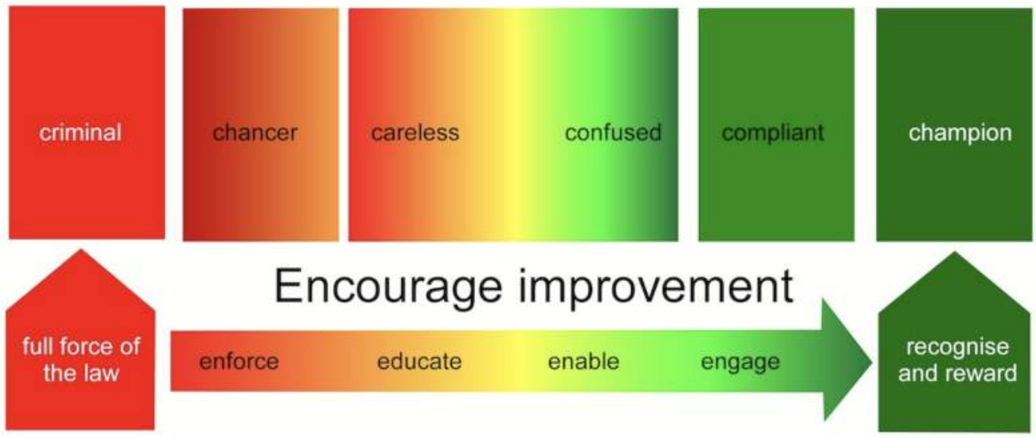

Figure 1.1: The responsive regulatory strategy as summarized in Braithwaite (2016) by the South Australian Environmental Protection Agency.
Inspectees who fall on the left side of the spectrum should be faced with severe enforcement. As one progresses to the right side of the figure, the enforcement strategy should change towards less coercive, reward-based strategies.

Because of the limitations of the typologies of compliance (see Section 1.2.1), Scholz (1984) notes that an integrative enforcement strategy produces better compliance outcomes than a single, fixed strategy like deterrence through fines or the threat of being caught. As a result, several researchers have attempted to integrate various typologies into a flexible enforcement strategy to provide an adaptable way for regulatory agencies to determine appropriate interventions. Notably, in the book *Responsive Regulation: Transcending the Deregulation Debate*, Ayres and Braithwaite (1992) propose a pyramid of enforcement strategies aiming to provide a flexible set of intervention strategies based on the type and nature of the offense. Known as responsive regulation, the essence of this pyramid is that the severity of the enforcement strategy should be aligned with the severity of the criminal behavior (see Figure 1.1). If the behavior of the inspectee improves, the severity of the enforcement can be reduced over time. This responsive regulatory framework formed the basis of subsequent studies on regulatory policy and the rise of risk-based regulation (see Baldwin and Black, 2008; Black and Baldwin, 2010; and Lodge and Wegrich, 2012).

In responsive regulation, the enforcement intervention is sensitive to the behavior of inspectees. Rather than being prescriptive, it leaves inspectors with the flexibility and freedom to develop targeted interventions suited to the particular offense. Responsive regulation led to the formation of several other intervention strategies, namely tit-for-tat, tripartism, self-regulation, and partial-industry intervention (Van der Heijden, 2020; see Table 1.2).

Intervention	**Strategy Description**
Tit-for-tat	The severity of the enforcement is aligned with the severity of the offense.
Tripartism	Empower other third-party organizations like other industry corporations or citizen associations to create an environment of cooperation.
Self-regulation	Encourage firms to set their own regulations that are publicly approved so that enforcement of the rules. can also be done publicly.
Partial-industry intervention	Target specific, competitive firms for inspection, with the idea that the natural competition of the market will also impact the portion of the market that is not inspected.

Table 1.2: Intervention strategies of responsive regulation (adapted from Van der Heijden, 2020).

While some studies prove that responsive regulation increases compliance (see Islam and McPhail, 2011; Christian, 2017; Zhu and Chertow, 2019), there remains several criticisms to responsive regulation. The main criticism is that it is effective only when inspectors and inspectees have multiple interactions over time (J. L. Short and Toffel, 2010; Van Duin et al., 2018). By nature of being responsive, the effectiveness of enforcement strategies can only be evaluated if there is additional capacity and resources for follow-up inspections to see if the intervention had the intended effect. Moreover, if inspectors are known not to escalate enforcement with repeated offenses, overall compliance is usually undermined (Van Erp, 2011).

ILT's Enforcement Strategy

At the ILT, a new enforcement strategy came into affect on May 14, 2022. Called the Landelijke Handhavingsstrategie (LHS), the National Enforcement Strategy, it is a flexible intervention strategy similar to Ayres and Braithwaite's responsive regulatory framework (see Section 1.2.1.1). The LHS aims to provide more transparency, predictability, and legal equality in the way in which inspectees are monitored (Inspectie Leefomgeving en Transport, 2022). This strategy evaluates the behavior of the offenders and its potential social damage to determine the appropriate intervention (see Figure 1.2), akin to the tit-for-tat responsive regulatory strategy (see Table 1.2).

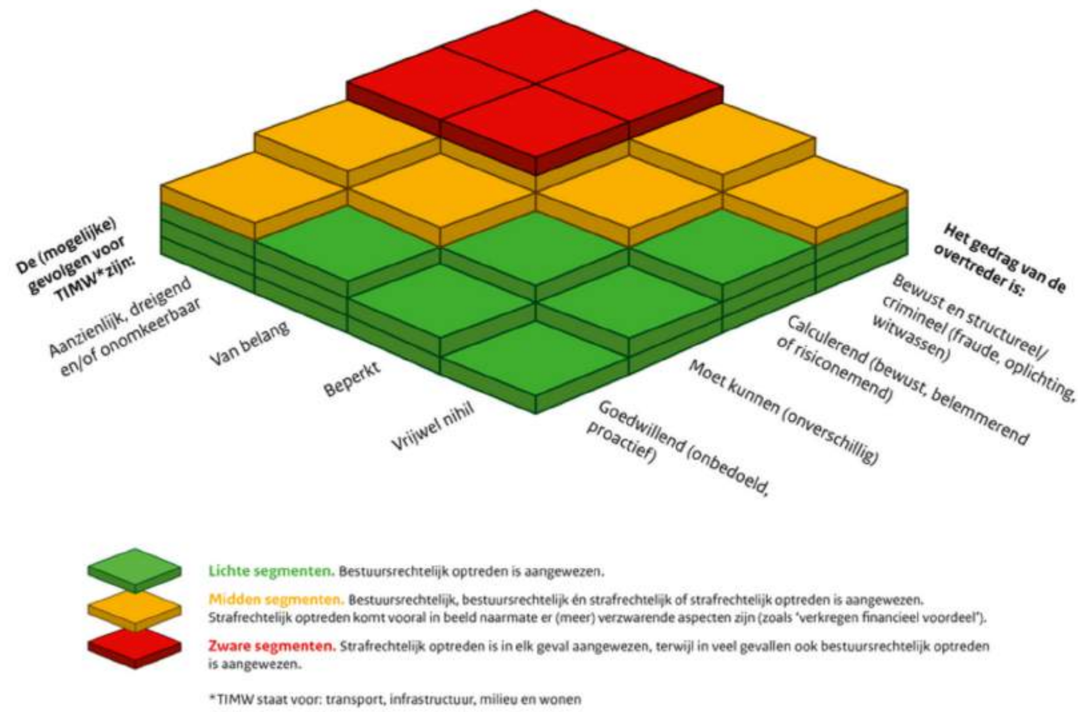

Figure 1.2: Landelijke Handhavingsstrategie (LHS) intervention matrix (Inspectie Leefomgeving en Transport, 2022).

The LHS shifts the way in which the ILT conducts inspections. Before the LHS, inspectors followed the "Interventieladder" (Intervention Ladder) policy where the inspectors were charged with choosing the standard, lowest level of enforcement; they only increased the enforcement severity after repeated offenses (Inspectie Leefomgeving en Transport, n.d.). Conversely, the LHS gives inspectors more freedom to choose from a menu of interventions at varying severity based on how inspectees behave. This strategy also shifts the purpose of inspections, encouraging inspectors to go beyond increasing compliance to focus on achieving the highest positive social impact (Inspectie Leefomgeving en Transport, 2022). The introduction of this new enforcement strategy warrants further study into its effectiveness, especially since there is a lack of knowledge on the causal drivers of inspectees' behavior.

1.2.2 Phenomenological Modeling

1.2.2.1 Phenomenological versus Mechanistic Modeling

Phenomenological models, also known as statistical models, represent observations of a particular target phenomena without postulating the underlying mechanism behind the phenomena (Rodrigue and Philippe, 2010; Frigg and Hartmann, 2020). In other words, a phenomenological model simply attempts to represent an observed relationship between variables that give rise to a particular phenomenon described by data (Hilborn and Mangel, 2013). A simple example of a phenomenological model given by Rodrigue and Philippe (2010) is the calendar system. Representing observations of daily and seasonal patterns, the calendar system has been used by many ancient cultures to accurately predict seasonal changes without necessarily understanding the physical causes of their observations. In fact, these cultures often explained the underlying mechanisms with divine narratives before physical mechanisms were discovered (Rodrigue and Philippe, 2010).

On the other hand, mechanistic models describe a relationship between two variables in terms of specific processes, rules, and equations where their outcomes found in data (Otto and Day, 2011; Hilborn and Mangel, 2013). In mechanistic models, parameters have their own unique definitions and can be measured independently from the outcomes found in data (Hilborn and Mangel, 2013; Connolly et al., 2017). Kendall et al. (1999) mention several strengths that mechanistic models offer. They explicitly include factors that are believed to influence the dynamics between variables and can predict outcomes based on these factors. Additionally, they provide a theoretical framework to model emerging behavior from different underlying mechanisms, providing further understanding of behavioral drivers within a system. The perfect rationality model (see Section 1.2.1) is mechanistic in nature, as it uses mathematical reasoning to determine an agent's behavior. In addition to the drawbacks of the perfect rationality model discussed in Section 1.2.1, mechanistic models also ignore the statistical properties of a system; in doing so, these models run the risk of making additional unrealistic assumptions that lead to "no sharp notion of goodness of fit" (Kendall et al., 1999). Phenomenological modeling allows the modeler to depict the relationship between variables without having to arbitrarily assign a mechanism or an equation to the relationship that may ignore the heterogeneity of agents and observations in the real world. Over time, the accumulation of phenomenological models can lead to a more accurate mechanistic discovery; one can argue that the accumulation of phenomenological calendar models led to the understanding of the laws of physics (Rodrigue and Philippe, 2010). Table 1.4 summarizes the distinction between phenomenological and mechanistic models.

Model Type	**Phenomenological Modeling**	**Mechanistic Modeling**
Aim	Describe and depict the relationship between two variables to extrapolate beyond meaured data and statistically predict future observations.	Depict why variables interact the way that they do and how outcomes are produced.
Object of model	Represents observed relationship between variables that give rise to phenomenon described in data.	Represents specific processes and rules that relate variables that result in outcomes found in data.
Parameters	Depends on statistical properties of the system.	Can be measured independently from the data and have their own unique definitions.
Risks & Limitations	Little insight into the underlying causes that cause empirical observations.	Ignores statistical properties of a system which may lead to unrealistic assumptions.

Table 1.3: Phenomenological versus mechanistic modeling.

1.2.2.2 Examples of Phenomenological Modeling

Phenomenological modeling has been widely used in the physical sciences, namely ecology, evolutionary biology, genetics, and fluid dynamics. Phenomenological modeling has provided scientists with a way to predict biological or physical processes and develop potential testable hypotheses that may explain the processes (White and Marshall, 2019). Yet, there are examples of phenomenological models that represent socio-technical systems, especially in the study of individual and interdependent decision-making.

McFadden (2001) pioneered the application of a phenomenological approach to social systems when he developed the discrete choice theory. This theory predicts people's behavior based on empirical observations. In his discrete choice model, McFadden (2001) uses an empirical measurement of individual utility derived from polling data. By using only observed data, the model accurately predicts an individual's choice to use the Bay Area Rapid Transit (BART) system in California, USA. In several applications that simulate commuting behavior, the model's predictions have a precision error of less than 2% (Contucci and Vernia, 2020). However, generalizations of the model to simulate other social issues show results that were incorrect and unrealistic (Contucci and Vernia, 2020). This led to the understanding that individual behavior is oftentimes dependent on the behavior of others within their social network. Brock and Durlauf (2001) build upon discrete choice theory by adding the effect of social interactions that cause individuals to conform their behavior to others. Their work provides insight into this observed phenomenon, showing that small changes in individual choices can lead to large changes in population behavior, particularly in non-cooperative settings. Further studies by Contucci and Ghirlanda (2007) and Barra et al. (2014) use similar phenomenological modeling techniques, layering the effects of social interactions onto discrete choice theory to predict immigration patterns that result from individual choices.

These studies led to the emergence of statistical physics, also known as the statistical mechanics approach to modeling social phenomena (Castellano et al., 2009). This approach is based on research that shows that small, local changes in behavior can influence collective behavior on a societal level which tends to be highly predictable (Castellano et al., 2009; Perc, 2019). Within the field of criminology, real-world observations show that crime tends to be concentrated spatially and temporally in urban areas (D'Orsogna and Perc, 2015). This understanding led to subsequent studies that model how crime stems from interactions with non-linear feedback loops, causing the observed phenomenon of crime clusters within complex social systems (see Brantingham and Brantingham, 1984; M. B. Short et al., 2008; Chainey and Ratcliffe, 2013; Alves, Ribeiro, Lenzi, et al., 2013; Alves, Ribeiro, and Mendes, 2013; and Picoli et al., 2014). Ball (2012) refers to this phenomenon as "hotspots" and attributes its emergence to social feedback loops. These studies show that it is possible to take empirical data and find statistical relations between variables that are precise enough to predict the evolution of macro-level behavioral patterns. This form of statistics-based, phenomenological modeling has been gaining recognition in the last decade; for example, the Santa Cruz police department in California uses empirical data and phenomenological models to predict future crime hotspots (Ball, 2012).

1.2.2.3 Advantages of Phenomenological Modeling

As research shifted from a classical economical approach towards an interactions-based understanding of individual behavior, phenomenological modeling provided a common framework to better comprehend complex socio-economic systems (Durlauf, 1999). Within the physical sciences, phenomenological models depict the correlation of behavior between one atom and its neighbor; similarly, phenomenological models can show how individual behavior relates to the behavior of other actors in the environment. This gives a common structure for discussing social problems by suggesting other ways to think about what drives social phenomena (Durlauf, 1999). Phenomenological modeling allows the modeler to represent the relationship between variables as observed in reality without running the risk of incorrectly incorporating mechanisms or causalities that do not exist in the real world. Moreover, a phenomenological approach can provide insight into how intervention strategies impact emerging collective behavior and the overall social system, providing a powerful tool for assessing the effectiveness of the strategies (D'Orsogna and Perc, 2015). Because phenomenological models are created from observations, their results are more easily translated to applications in the real world (Van der Schaaf, 2019).

1.2.2.4 Limitations of Phenomenological Modeling

By nature of being based on observed data, phenomenological models may not be applicable to other environments outside of the one in which data was collected (Van der Schaaf, 2019). Phenomenological models require empirical data without which the model cannot be created. Additionally, because phenomenological models are based on historical observations, their statistical inferences may not encompass the wide range of individual choices. Unlike atomic particles, humans can learn and adapt; their past behavior does not necessarily dictate their future behavior (Fortunato et al., 2013). They may change their mind in the future or act like they never have before. Phenomenological models may also be unstable, meaning that a small change in inputs yields drastically different results (Van der Schaaf, 2019). While phenomenological models may successfully find statistical correlations between variables, it does not prove causation (Fortunato et al., 2013). They are intended purely to model the statistical relationships from observed phenomena and do not prove the underlying causes behind how the system functions. Thus, multiple models may fit empirical observations without being incorrect; it is crucial that phenomenological models are validated adequately to the context at hand (Van der Schaaf, 2019). To reduce model instability, multiple phenomenological models that fit the data can be aggregated into an averaged model. Yet, the averaged model must also be validated to ensure that it fits with empirical observations (Van der Schaaf, 2019).

While there may be commonalities between physical and social systems as discussed in the preceding sections, there are crucial differences that make modeling social systems much more complex (Ball, 2012; Fortunato et al., 2013; Van der Schaaf, 2019). Modeling observed data can lead to an endless number of parameters that could potentially describe the system, and it is unreasonable to gather measurements for all of them (Van der Schaaf, 2019). Moreover, these parameters may be challenging or even impossible to quantify (Ball, 2012). Additionally, human behavior can change over time and can have multiple dimensions spanning economics, ethics, and culture (Ball, 2012; Perc, 2019). Finding the appropriate metric remains a challenge in complex social systems, especially with the limited capacity of inspectors.

On a larger scale, social systems tend not to follow a set of rules. While circumstances may produce a particular observed phenomenon in one situation, it may produce a different phenomenon in another situation (Ball, 2012). Outcomes in real social systems may be dependent on a wide range of contingencies that produce too much variation in the system, making it tough to obtain repeatable results from a phenomenological model (Ball, 2012). Again, multiple models with completely different assumptions could still fit the same empirical phenomenon which makes it tough to pinpoint the exact influencing factors impacting the system (Ball, 2012; Van der Schaaf, 2019).

Even with these limitations and complexities, phenomenological modeling is still possible provided there is adequate empirical data. A pure phenomenological approach requires a representative dataset of at least two measured variables to find correlations between them. Without sufficient data, models will have to rely on theories at which point it is no longer purely phenomenological.

1.2.3 Agent-based Modeling

Agent-based modeling (ABM) is a modeling method that consists of multiple agents that have various characteristics and behaviors (Edmonds and Gershenson, 2015). These agents are independent and autonomous; they interact with one another and their environment to decide on their next action (Macal and North, 2005). Throughout the simulation, they may exercise their agency and learn from their past decisions over time (Edmonds and Gershenson, 2015). This section discusses ABM's potential for investigating inspection and enforcement strategies while discussing its purpose, advantages, limitations, and examples of its application within law enforcement.

1.2.3.1 Purposes and Advantages of Agent-based Modeling

As a reaction to the criticisms of the perfect rationality model (see Section 1.2.1), Simon (1957) presented the concept of bounded rationality. He proposed a more realistic depiction of the decision-making process by adjusting the perfect rationality model to include two new features: 1) a lack of access to perfect information and 2) the finite nature of an individual's capacity to process it. Moreover, Simon (1957) points out that decision-making is a resource- and time-consuming activity. Therefore, agents "satisfice"; they make the best possible choice with the information and resources they have (Simon, 1957). While the bounded rationality concept applies to inspectees and whether they choose to comply or violate, it also applies to inspectorates. Bounded rationality is at the core of an inspectorate's action dilemma. Inspectorates conduct law enforcement in an uncertain environment where they have incomplete information and a limited amount of resources to process available data (De Bruijn et al., 2007). Yet, they are mandated to find an acceptable compliance outcome in an industry. Given the "satisficing" principle associated with bounded rationality, inspectorates may not be able to make the absolute, optimal choice; this can only happen with perfect rationality, which is not realistic (see Section 1.2.1). However, inspectorates still need to make decisions that meet their compliance targets within an industry, given limited information about the sector and a finite number of resources to test their inspection and enforcement strategies.

According to Marin et al. (2020), ABMs have the ability to model bounded rationality within its agents. It has the potential to model agents without presuming their underlying motivations, depending on how the model is conceptualized. ABMs create agents and specify the interactions between them so that the macro-level consequences of their interactions can be studied (Gräbner, 2016). In this way, it provides the ability for the modeler to incorporate bounded rationality in the agents while embedding them in an environment (Marin et al., 2020). The agents work with the information available to them and make decisions to the best of their ability. As the agents interact with the environment and each other, the system as a whole also evolves, thereby requiring agents to readapt to the new environment shaped by their previous interactions.

ABMs have several advantages and applications for studying human behavior and developing effective strategies. First, they are a useful tool in the policy development process, as they can adequately represent an observed or theoretical phenomenon (Edmonds and Gershenson, 2015). By describing a particular phenomenon, ABMs can help to explain possible theories, explore different scenarios, and predict and weigh future outcomes (Epstein, 2008; Edmonds and Gershenson, 2015; Edmonds, 2017a). For "dynamic and complex phenomena" which Edmonds (2017a) defines as situations where several mechanisms interact over time, ABMs can offer a "direct representation without theoretical restrictions." This direct representation ensures there is consistency between the represented entities and their interactions.

In addition, ABMs are highly flexible and adaptable; changes to agent behavior can be made easily according to expert knowledge and random effects can be easily included (Barbaro, 2015; Edmonds and Gershenson, 2015). This adaptability allows more complexity to be captured compared to conventional mathematical models (Barbaro, 2015). ABMs can provide valuable insight into complex problems and help policymakers better understand complexity so that they devote resources more effectively and efficiently (Edmonds and Gershenson, 2015; Edmonds, 2017b). According to Edmonds and Gershenson (2015), ABMs can be a part of a decision-making process to decide on what adaptive strategies to take. They can serve as a base for subsequent generalizations and scenario development (Edmonds, 2017a). Moreover, ABMs can illuminate the dynamics of inspectee behavior, helping to show trends in behavior and serving as a 'virtual laboratory' to test out various intervention strategies before inspectorates deploy them in real life (Van Asselt et al., 2016).

1.2.3.2 Examples of Agent-based Models in Literature

In existing literature, there are examples of ABMs that simulate illicit criminal markets and explore the effectiveness of intervention strategies. Hoffer et al. (2009) employ an ABM to examine the heroin market in Denver, Colorado in the 1990s. They use observational data and link it to agent behaviors to see how it impacts the collective local market. Though their model simulation is experimental, the ABM proved to be a novel method that helped the experimenters better understand the dynamics of the heroin market and the outcomes they produce. Jones et al. (2010) use ABMs to explore the effectiveness of two police deployment strategies: increased patrolling throughout the crime hotspots and barring criminals from entering the crime hotspots. Their ABM shows that the deployment of law enforcement agents successfully decreases the crime rate, though its reduction depends on how many police agents are deployed and distributed. Patrolling the crime hotspot requires a smaller number of police agents than barring criminals at the periphery. Hegemann et al. (2011) use an ABM to simulate the interactions between rival street gangs using geographic information and gang data. They conclude that their ABM is a flexible enough model to test potential hypotheses of gang dynamics and formation, though it does not identify the particular causes of the formation of gangs. The simulation accurately represents the social phenomena of rival gang interactions; therefore, they recommend that intervention strategies can be incorporated and tested with the model.

In addition, there are examples of ABMs that simulate the inspection environment, though these are few. Van Asselt et al. (2016) use an ABM to simulate compliance behavior and evaluate the effect of various inspection strategies in a study of pig farmers using antibiotics. They incorporate social, economic, cognitive, and institutional factors into the behavior of pig farmers. Their model shows that the biggest factor for compliance is social pressure; if pig farmers are socially accepted, they positively influence each other which increases compliance rates. They also simulate various intervention scenarios which informs inspectorates of where to most efficiently allocate resources. Van der Voort et al. (2020) use an ABM to determine how inspectorates should use data-driven risk-based inspections without introducing long-term negative bias in data. They conclude that the best strategy is to maintain a certain percentage of random inspections while conducting data-driven inspections; this produces the best data quality over time despite inspectorates propensity to overestimate compliance rates. Smojver (2012) model the interaction between banks and bank supervisors using an ABM to alert bank supervisors to the rules that banks are most likely to violate. Knowledge of these risk areas led supervisors to adapt their intervention strategy to decrease the number of violations. Korobow et al. (2007), Hokamp and Pickhardt (2010), Llacer et al. (2013), and Andrei et al. (2014) create an ABM to model tax compliance within social networks, applying them to different cases. Andrei et al. (2014) conclude that an agent's network has a significant impact on tax compliance, especially in a centralized network with large fines. In the Spanish case, Llacer et al. (2013) find that social pressure improves compliance only in low to medium deterrence conditions. Korobow et al. (2007) find that aggregate compliance is higher when inspectees do not know the payoffs that are available to their neighbors. Hokamp and Pickhardt (2010) discover that ethical behavior patterns and lapse time effects significantly decrease tax evasion compared to audit probability or the tax rate.

1.2.3.3 Limitations of Agent-based Modeling

To model complex systems using ABM, system components must be examined in relation to one other; interactions must be studied to see if they generate "novel information" from internal system processes (Melchior et al., 2019). Doing so inevitably leads to abstractions of reality (Edmonds and Gershenson, 2015). Abstractions are needed to simplify the complexity of the system to zoom in on the problem at hand. Simplification requires choices to be made by the modeler. The modeler must critically assess the validity of these choices and weigh the trade-offs between choices. This assessment is necessary as a modeler's choices can change the model's fit for purpose for studying its context. For example, simplifying a model may lead to less validity in the real world; models with more abstractions may have a quality of generality at the expense of formality (Edmonds and Gershenson, 2015). To create a useful model of a complex environment, there must be a clear purpose and context for the model. By nature of being complex, the model may produce new processes and unexpected results (Edmonds and Gershenson, 2015). Clarity of focus ensures that the model's results are not over-extrapolated beyond its intended purpose and context. For policymaking, a clear purpose ensures that people do not "fool themselves" by relying too much on the model and hiding behind its opacity (Edmonds and Gershenson, 2015). If models become too opaque, it may be tough to connect it to the real world which diminishes the value of the model for policymaking (Melchior et al., 2019).

1.2.4 Phenomenological Approach to Agent-based Modeling

The network of enforcement is a complex system containing autonomous agents that interact with one another. Oftentimes, observed phenomena are noticed by inspectors before the underlying motivations are understood or studied systematically. Phenomenological modeling is a method that can analyze behavioral phenomena without presupposing the multitude of potential underlying motivations that may drive an individual's decision to act lawfully or criminally. It can investigate behavioral dynamics beyond existing behavioral theories of compliance such as peer pressure effects, thereby avoiding the drawbacks of the perfect rationality model (Durlauf, 1999; see Section 1.2.1). Historically, scientists have used ABM to understand various complex systems function, showing how interactions give rise to system behavior (Ball, 2012). While an agent-based approach can help to test the dynamics of behavior in an complex and interdependent social system while using a bounded rationality depiction of agents (see Section 1.2.3.1), it has been rarely used in the policy development process (Melchior et al., 2019). Combining a phenomenological approach with ABM provides a way of investigating complexity starting from observations rather than theory. ABMs are a natural expansion of phenomenological modeling of behavior (Ball, 2012), yet examples of a phenomenological approach to ABM for investigating the inspection environment are scarce. Table 1.4 summarizes the differences between a traditional approach and phenomenological approach to ABM.

Approach	Traditional Approach to ABM	Phenomenological Approach to ABM
Starting Point	Bottom-up; start from simple atoms that follow deterministic rules.	Top-down; start from identification of a macro-level phenomenon.
Method	Apply theories and concepts to the target domain by calibrating micro-interactions with measured, independent parameters.	Use statistical patterns to calibrate micro-interactions of the observed phenomenon. Assumptions are determined empirically to fit the observed data.
Purpose	Show how micro-interactions generate and lead to the emergence of macro-level behavior over time.	Reproduce existing empirical observational data to simulate possible future scenarios. If applicable to the study at hand, generate possible micro-interactions that led to the observed, macro-level phenomenon.

Table 1.4: Phenomenological versus traditional approach to ABM (adapted from Schinckus, 2019).

For the ILT, a phenomenological approach to ABM provides several benefits. First, it offers a tool for them to explore unknown aspects of the inspection environment using statistical analysis of historical data. Rather than basing decisions on models that rely exclusively on behavioral theories that may not be realistic, a phenomenological approach allows inspectorates to incorporate data into the policy development process. Where there is little insight into the motivations of inspectees or a lack of empirical data, the logic of ABMs can fill in the gaps (see Section 1.2.3.1). Using data to build the model gives the ILT more certainty that the model depicts real-world behaviors.

More importantly, a phenomenological approach to ABM strengthens the ILT's ability to address their action dilemma. The ILT has incomplete information about inspectees and their decision-making process. Moreover, they have limited resources to analyze existing data. Yet, they must continue to try to influence compliance behavior in an industry. With a phenomenological approach to ABM, the ILT can reproduce empirical data to explore the impact of their interventions on compliance outcomes. Then, intervention strategies can be tested to see how they function within the inspection environment. Specifically, the ILT's new enforcement strategy (the LHS; see Section 1.2.1.1) can be evaluated along with known characteristics of inspectees' behavior that have been observed and discovered by statistical analysis. A phenomenological ABM helps inspectorates pinpoint where and how to deploy their resources to produce the best outcomes given the diversity of observed inspectee behaviors.

Chapter 2

Methodology

With few existing studies that use a phenomenological ABM to model the inspection environment, an exploratory research and modeling approach are used to apply a phenomenological approach to ABM. The methods used to address each research sub-question are shown in Table 2.1.

1	**What are the theoretical foundations for the application of phenomenological and agent-based modeling techniques for studying compliance behavior?**
	Method: Literature review (Section 1.2).
2	**What theories underpin the conceptualization of an ABM depicting the inspection environment?**
	Method: Review of the existing ABM created by Knol (2021) for the ILT and the underlying theories that informed the model (Section 2.1.1).
3	**What data is needed to identify observed behavioral phenomena in the inspectee population?**
	Methods: • Unstructured interviews and informational meetings with ILT experts to: – identify any behavioral phenomena that they notice during their inspections (Section 2.1.2); and – request any data that has been collected from inspections. • Data analysis of available inspection data to identify other behavioral phenomena in addition to those discovered during unstructured interviews and informational meetings (Section 2.1.2). • Participatory observation of ILT activities, including weekly meetings with ILT experts to validate assumptions where there is missing data.
4	**How can behavioral phenomena be conceptualized into an ABM?**
	Method: Development of a conceptual model of the phenomenological ABM in Netlogo (Chapters 3 and 4).
5	**How can the effectiveness of inspection and enforcement strategies be investigated under different scenarios of the inspectee population?**
	Method: Simulation of scenarios and analysis of model results (Chapter 5).

Table 2.1: Methods for each research sub-question.

Note that the observed behavioral phenomena identified as part of sub-question #3 refers to statistical patterns of behavior found in data. It also includes qualitative data which are inspectors' personal observations of inspectees' behavior as they conduct inspections.

2.1 Theoretical and Empirical Background

2.1.1 Inspection Game Theory

Knol (2021) created an ABM for the ILT based on inspection game theory. Inspection games are a game theoretic mathematical model that simulates interactions between two types of actors: inspectors and inspectees. The first inspection game model was developed in the 1960s by the United States Arm Control and Disarmament Agency (USACDA) to study arms control inspections performed by the International Atomic Emergency Agency (IAEA) (Avenhaus et al., 2002). In addition, inspection games were used in the study of tax inspections (see Klages, 1968 and Borch, 1982) and environmental pollution problems (see Bird and Kortanek, 1974; Höpfinger, 1979; Brams and Kilgour, 1988; Weissing and Ostrom, 1991; and Güth and Pethig, 1992). In an inspection game, inspectors check to see if the inspectees comply with legal regulations (Avenhaus et al., 2002). Inspectors want to thwart illegal activity of inspectees and detect violations, while inspectees choose whether to comply or not. In a two-player game containing one inspector and one inspectee, the decision matrix is given by Table 2.2 where the payoff of the inspector and inspectee are represented by o and e, respectively.

		Inspector	
		Inspect (p)	Not Inspect (1-p)
Inspectee	Comply (q)	$o_{c,i}$ $e_{c,i}$	$o_{c,ni}$ $e_{c,ni}$
	Not Comply (1-q)	$o_{nc,i}$ $e_{nc,i}$	$o_{nc,ni}$ $e_{nc,ni}$

Table 2.2: Utility values of choices in a two-player inspection game (adapted from Knol, 2021).

Each player in the inspection game is subject to various costs while they try to maximize their utility. Inspectors are bound by resource and time constraints; therefore, they want to avoid inspecting actors they know to be compliant. To maximize their utility (e) when an inspectee is compliant, they would opt not to inspect. Conversely, when an inspectee is not compliant, choosing to inspect would increase the inspector's utility. Similarly, inspectees would choose to maximize their utility (o) by complying when an inspection is certain and not complying when they know they will not be inspected.

When the players act in this way, there is no Nash equilibrium, the point at which no player has the incentive to deviate from their decision (Knol, 2021). Player A will always choose the strategy that gives them the highest utility in response to Player B, which causes Player B to adjust their strategy to the one that gives them the highest utility. As long as the players maintain the same decision (in other words, maintain a pure strategy), a Nash equilibrium solution is not possible.

However, a mixed strategy equilibrium is possible as long as the players use a probabilistic strategy where the inspector and inspectee calculates the probability they will inspect and comply, respectively. The mixed strategy equilibrium is given by $p*$ and $q*$ in Equation 2.1 for the inspector and inspectee, respectively, where p is the probability the inspector will inspect and q is the probability the inspectee will comply (Knol, 2021).

$$p* = \frac{e_{c,ni} - e_{nc,ni}}{e_{c,ni} - e_{nc,ni} + e_{nc,i} - e_{c,i}} \tag{2.1a}$$

$$q* = \frac{o_{nc,ni} - o_{nc,i}}{o_{nc,ni} - o_{nc,i} + o_{c,i} - o_{c,ni}} \tag{2.1b}$$

Previous studies of the mixed strategy equilibrium show counter intuitive results that call into question the validity of inspection game results (Knol, 2021). For example, when fines are increased, the probability that the inspectee will comply ($q*$) increases as expected. However, it simultaneously decreases the probability that the inspector will inspect ($p*$), which is a less intuitive outcome. Moreover, an increase in the cost of compliance results in a higher probability of inspectors choosing to inspect ($p*$).

In addition to these counter-intuitive results, the inspection game model has other limitations. As shown in Equation 2.1, the players must have knowledge of the utility functions of the other players. In this way, the inspection game model assumes the players are perfectly rational when in reality, they do not have access to complete information (see Section 1.2.1). It is already difficult to estimate the utility function of a single player, let alone other players (LaValle, 2006). This information is oftentimes unavailable to other players in the inspection environment. Moreover, it is tough to model a randomized strategy where even rational actors can display unexpected behavior (LaValle, 2006). Past compliance behavior is not always an adequate proxy of future compliance as inspectees can change their willingness to comply abruptly and irrationally (Chalfin and McCrary, 2017). Game theory assumes players are purely utility-maximizing and perfectly rational; it does not account for the fact that players could choose to act at their own expense if given another incentive. Knol's ABM assumes that all agents are economically-driven, which limits the diversity of actor motivations (see Section 1.2.1). For example, it assumes that higher fines result in a lesser risk of violations though in reality, this is not always true for all inspectees. Moreover, the model does not consider that behavior is contagious, and individual decisions are significantly driven by group pressures (Durlauf, 1999; Keizer et al., 2008). Knol's ABM does not incorporate social interactions between inspectees that may influence how they choose to behave. Furthermore, it does not evaluate the effectiveness of responsive enforcement strategies (see Section 1.2.1.1) nor does it incorporate statistical properties of the inspection environment. Therefore, the risk of making unrealistic assumptions that lead to "no sharp notion of goodness of fit" remains (Kendall et al., 1999).

2.1.2 Empirical Observations of Behavioral Phenomena

A data-driven, phenomenological approach can bypass the limitations of inspection game theory. Phenomenological modeling allows the modeler to depict the relationship between variables (e.g., compliance rates and number of inspections) without having to arbitrarily assign a mechanism to the relationship that may ignore the heterogeneity of agents and observations in the real world. A phenomenological approach to ABM starts by identifying macro-level behavioral phenomena (see Table 1.4). To do so, analysis of both quantitative and qualitative data must be conducted. During unstructured interviews with ILT's experts, three categories of behavioral phenomena describing how inspectees decide whether to comply or violate were identified (see Table 2.3).

#	Behavior based on:	Description
1	Peer pressure	Inspectees conform to the behavior of their neighbors.
2	Inspections	Inspectee either: · Become more compliant with increasing inspections, · Become more non-compliant with increasing inspections, or · Do not change their behavior with increasing inspections
3	Enforcement	If inspectees can absorb the severity of enforcement, they will continue to violate. Otherwise, they become compliant.

Table 2.3: Behavioral phenomena studied in this research.

Inspectees' behavioral response to peer pressure (#1 in Table 2.3) and enforcement (#3 in Table 2.3) were identified by inspectors' personal observations, as there is no quantitative data on the network effects of the Dutch inland shipping sector or inspectees' reaction to enforcement interventions. Nevertheless, the ILT has access to the Inspectieview Binnenvaart database for the inland shipping sector (Binnenvaart in Dutch) where there is statistical evidence that inspectees' behavior is correlated with inspections (#2 in Table 2.3). This database has records of inspections conducted between March 2, 2015 and March 16, 2020 from five agencies: 1) ILT, 2) Ministry of Infrastructure and the Environment (Rijkswaterstaat), 3) Dutch National Police, 4) Port Authority of Rotterdam, and 5) Port of Amsterdam. A summary of the Inspectieview Binnenvaart data and the breakdown of ships by average violations are shown in Table 2.4 and Figure 2.1, respectively.

Inspectieview Binnenvaart Data	
Total inspections conducted over 6 years	40,690
Total violations found in 6 years	17,991
Average compliance rate over 6 years	55.8%
Total inspections ILT conducted in 6 years	6,565
Average number of inspections ILT conducted per year	1,313
Total unique ships	7,878

Table 2.4: Inspectieview Binnenvaart dataset summary.

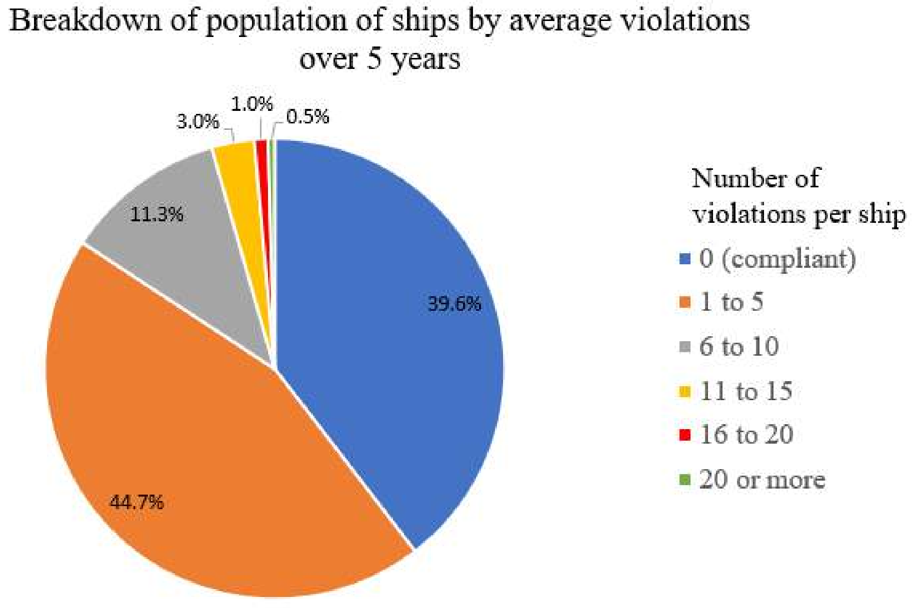

Figure 2.1: Average violations per ship
in the population of 7,878 inland ships (adapted from Meester, 2021).
On average, approximately 40% of inspected ships are compliant, with a majority of inspectees having less than 10 violations.

In a study by Meester (2021), group-based trajectory modeling (GBTM) is used to identify the development trajectories of subgroups of inland shipping inspectees. GBTM uses a maximum likelihood estimation rather than a cluster analysis to find clusters of actors with similar trajectories (Nagin, 2010). It is a special type of finite mixture modeling that finds latent groups within a population (Nagin, 2010). Therefore, GBTM is useful for discovering subgroups of inspectees that behave differently and has the capacity to analyze longitudinal data of multiple groups (Meester, 2021).

The results of Meester's GBTM analysis of the Inspectieview Binnenvaart dataset show that there are four main subgroups of inland shipping inspectees (Figure 2.2). The "descenders" and "strong descenders" are inspectees whose compliance rate improves with increasing inspectinos over time. In other words, these are deescalatory inspectees as they deescalate the severity of their offense with each inspection. The "increasers" are those who do the opposite; they are escalatory inspectees who increase the severity of their offense with each inspection. The "no change" inspectees are those whose behavior does not change with inspections. This result informs behavioral phenomena #2 in Table 2.3.

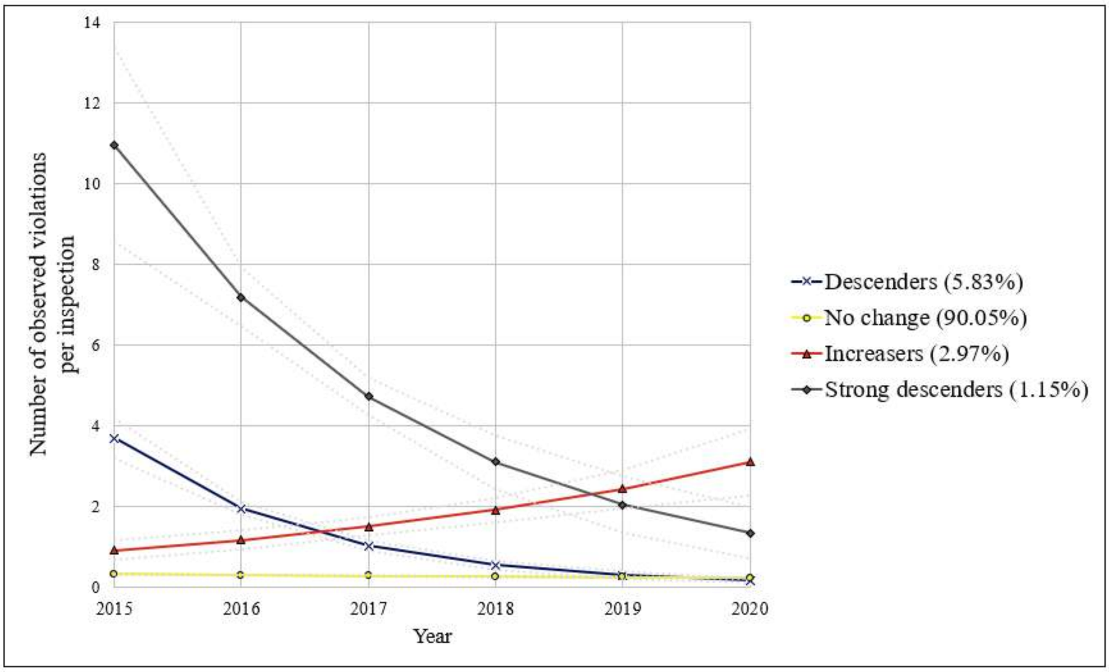

Ship-level patterns based on the number of observed violations per inspection over the 2015-2020 period (n=7757).

Figure 2.2: Ship-level compliance patterns
(Meester, 2021; translated from Dutch in Google Translate).

Approximately 7%, 3%, and 90% of inspectees are deescalatory, escalatory, and "no change" inspectees, respectively (Figure 2.2). Table 2.5 shows the average violations per inspection for each of these inspectee subgroups.

Inspectee Subgroup	% of population	Avg Violations per Inspection
Descenders (deescalatory)	5.83%	5.80
Strong Descenders (deescalatory)	1.15%	8.47
Increasers (escalatory)	2.98%	5.62
No change	90.05%	1.8

Table 2.5: Average violations per inland shipping inspection by inspectee subgroup (adapted from Meester, 2021).

The breakdown of non-compliant inspectees by type of offense is shown in Table 2.6.

Offense Type	%
Minor Violations	18.4%
Medium Violations	67.6%
Serious Violations	2.5%
Unclassified Violations	11.5%

(a) Including unclassified violations

Offense Type	%
Minor Violations	20.8%
Medium Violations	76.4%
Serious Violations	2.8%

(b) Excluding unclassified violations

Table 2.6: Breakdown of inland ships by offense type (adapted from Meester, 2021).

The average number of violations by offense type over the six year period between 2015 and 2022 is shown in Table 2.7.

Subgroup / Violation Type	Descenders	Strong Descenders	Increasers	No Change	Total Avg Violations
Minor	0.99	0.52	1.56	0.92	3.99
Medium	4.31	6.97	2.71	0.68	14.67
Serious	0.16	0.05	0.19	0.14	0.54
Unclassified	0.34	0.94	1.15	0.07	2.5
Total Avg Violations	5.8	8.48	5.61	1.81	21.7

Table 2.7: Average violations of each inland shipping inspectee subgroup between March 2, 2015 and March 16, 2020 (adapted from Meester, 2021).

It is worth noting that the differences between minor, medium, and serious violations were not defined in the Inspectieview Binnenvaart dataset. Each agency has their own method of categorizing violations which means that, for example, what one agency deems as a minor violation may be a medium level violation for another agency. The results in Tables 2.6 and 2.7 reflect the three tiers of violation types that are denoted in the data.

Table 2.8 shows the breakdown of enforcement interventions given to non-compliant inspectees in each inspectee subgroup.

	Descenders	Strong Descenders	Increasers	No change
Warning	69%	77%	57%	56%
Administrative	12%	10%	10%	6%
Police Report	4%	1%	3%	8%
Suspension or Shutdown	2%	0%	0%	1%
Unclassified	14%	12%	31%	30%

Table 2.8: Enforcement interventions for each inland shipping inspectee subgroup (adapted from Meester, 2021).

A majority of the enforcement interventions are warnings, specifically 67.6%. Administrative interventions account for 10.4% of total interventions. Police reporting and suspension or shutdown represent 2.7% and 0.5% of interventions, respectively. The remaining 18.7% are unclassified interventions.

2.1.2.1 Limitations of the Inspectieview Binnenvaart Data Analysis

Because the Inspectieview Binnenvaart dataset comes from five agencies, the data was highly aggregated, leading to some loss of data. Each of the agencies collects data in their own way and have distinct procedures for conducting inspections. Therefore, combining data from five agencies into one large dataset requires assumptions on how data labels from one agency correspond with those of another. Notwithstanding, an aggregated dataset offers a wide range of information that includes foreign and domestic ships that operate in the Netherlands; this provides a more comprehensive view of the sector. Yet, the Inspectieview Binnenvaart dataset does not define data labels which could have different definitions in each agency. As mentioned previously, the categorization of minor, medium, and serious violations is subjective and unique to each agency; the dataset does not define the differentiation between the violation types. In addition, operational dissimilarities between the agencies cause discrepancies in the definitions of variables. For example, each inspection agency has different inspection strategies, and the data does not indicate how inspectees are targeted for inspections. The inspection data provided by the ILT is biased according to how inspectors subjectively decide which ships to inspect. At the time the data was collected, ILT's inspectors followed the "Interventieladder" policy (see Section 1.2.1.1). Moreover, each agency inspects for violations of different regulations, so the type of inspection and its thoroughness may vary between agencies. The enforcement interventions that are executed by each agency also differ, as each agency has its own mandate.

Despite these limitations, the GBTM methodology is a useful tool to find behavioral trajectories of different subgroups within the inspectee population. However, it is limited to the study of one type of trajectory over time. Meester's 2021 study characterizes the trajectory of violations only in relation to inspections and not any other factors (e.g., type of violation). In other words, the GBTM analysis assumes that other factors are not time sensitive, which limits the understanding of how other external factors impact the trajectory of violations over time.

Chapter 3

Conceptualization of the Phenomenological ABM for Inspections (PABMI)

This chapter details the conceptualization of the Phenomenological Agent-based Model for Inspections (PABMI) using the theoretical and empirical background from Section 2.1. This section details the construction of the ABM after identifying three categories of observed behavioral phenomena (see Table 2.3 and Section 2.1.2).

3.1 PABMI: Model Purpose

The purpose of the PABMI is to identify which inspection and enforcement strategy is the most effective at increasing compliance given the three categories of behavioral phenomena exhibited by inspectees (see Table 2.3). In addition, the PABMI is intended to spur conversation within the ILT about the intricacies of the inspection environment that may not yet be known.

3.2 PABMI: Model World Setup

The PABMI is developed in Netlogo version 6.2.2 (Willensky, 2021) on Windows 11 Home. Developed by Willensky (2021), Netlogo is an openly available software package to create ABMs. Netlogo is the chosen software because the code is easy to understand, and it is simple to use (Van Dam et al., 2012). In addition, it has a graphical user interface that provides a visualization of the model. This provides a crucial benefit for demonstrating the model's capabilities to inspectors at the ILT. The interface allows inspectors to visualize the evolving simulation and adjust model parameters easily, making it more practical and easy to use.

The PABMI is set up with one agent type – inspectees – with a user-definable population size (*n-inspectees*). These inspectees are agents that are created and placed in random locations in the model world. Each inspectee agent has 15 state variables (Table 3.1).

#	State Variable	Description
1	*compliance*	Compliance level ("compliant", "unintentional", "conscious", or "criminal")
2	*offense-severity*	The severity of the offense
3	*n-historical-violations*	Number of historical violations
4	*n-unintentional-offenses*	Number of unintentional offenses
5	*n-conscious-offenses*	Number of conscious offenses
6	*n-criminal-offenses*	Number of criminal offenses
7	*compliance-rate%*	Compliance rate (%) (defined as the n-compliant / n-times-inspected)
8	*n-compliant*	Number of times the inspectee complied
9	*n-noncompliant*	Number of total violations when inspected
10	*n-times-inspected*	Number of times the inspectee was inspected
11	*inspection-candidate*	Indicator of whether the inspectee was chosen for inspection (boolean true or false)
12	*inspection-reaction*	Inspectee's reaction to inspection ("escalate" or "deescalate")
13	*enforcement-severity*	Indicator of the severity of enforcement performed on non-compliant inspectees by inspectors (0-1, where 1 is the most severe)
14	*absorbance-capacity*	Indicator of the inspectee's capacity to absorb the severity of enforcement
15	*xcor* and *ycor*	X and Y location coordinates of the inspectee agent

Table 3.1: State variables of the inspectee agents.

Upon setup, the inspectee agents are assigned values for four state variables which are their initial characteristics: 1) Number of historical violations, 2) Compliance status, 3) Severity of offense, and 4) Absorbance capacity indicator. The PABMI was designed so that user-definable parameters can be easily changed with sliders on the Netlogo interface. This enhances the flexibility of the model, as parameters can be adjusted to reflect different population characteristics.

1. Number of historical violations

The number of historical violations (*n-historical-violations*) is the agent's record of past violations. This state variable is determined phenomenologically as it is based on empirical data (see Section 2.1.2). The average number of violations by inspectee subgroup over the six-year time period of the Inspectieview Binnenvaart data is used for this initial condition (see Table 2.7).

2. Compliance level

The compliance level (*compliance*) of each inspectee agent can be one of four options: "compliant," "unintentional," "conscious," and "criminal." As defined by the LHS (see Figure 1.2), unintentional offenders are indifferent, benevolent, and potentially incompetent actors whose violation impact is limited to none. Conscious offenders are calculative, risk-taking actors whose violation has a limited to moderate impact. Criminal offenders violate in a structural and coordinated way (e.g., money laundering and fraud) with their violations having significant, threatening, and/or irreversible impact.

The initial share of inspectees at each of the four compliance levels is determined phenomenologically from the Inspectieview Binnenvaart dataset. The initial share of compliant inspectees (the *percent-compliance* parameter) is set to 39.6% (Figure 2.1); this means that 60.4% of the population is non-compliant at the start of the simulation. Of the non-compliant inspectees, 20.8%, 76.4%, and 2.8% of them are defined to be unintentional, conscious, and criminal offenders, respectively (see Table 2.6b). These percentages define the *percent-unintentional*, *percent-conscious*, and *percent-criminal* parameters that represent the initial share of inspectees that are unintentional, conscious, and criminal offenders, respectively.

Note that while the Inspectieview Binnenvaart dataset labels violation types as "minor," "medium," and "serious" (Table 2.6), the LHS uses "unintentional," "conscious," and "criminal" to distinguish between violation types (see Figure 1.2). Because the LHS was adopted after data was collected for the Inspectieview Binnenvaart database, the labels used to categorize different violation types do not match. Moreover, the LHS is specific to the ILT while the categorization of "minor," "medium," and "serious" violation types are used for inspections across the five inspection agencies. For this research, "minor," "medium," and "serious" violations are assumed to be "unintentional," "conscious," and "criminal" levels of compliance, respectively. This assumption is based on the definition in the LHS that unintentional, conscious, and criminal offenders have the lowest, medium, and highest offense severity, respectively. While this categorization may not be valid in the real world, the lack of definition of these violation types requires this assumption to be made for the model conceptualization.

3. Severity of offense

To indicate the severity of the violation, each inspectee subgroup is assigned an *offense-severity* indicator based on their compliance level (Table 3.2).

Compliance Level	Offense Severity
Compliant Inspectees	0
Unintentional Offenders	10
Conscious Offenders	20
Criminal Offenders	30

Table 3.2: Offense severity index for each compliance status.

This assignment of the offense severity by compliance level is based on the ILT LHS framework derived from existing behavioral theories; the LHS framework was not developed empirically, but rather theoretically on the basis of the theory of responsive regulation (see Secion 1.2.1.1). Nevertheless, the LHS's categorization of inspectees based on the severity of their offense is necessary for the empirical analysis of the inspectee behavior. As previously mentioned, the Inspectiview Binnenvaart data categorizes violation types as "minor," "medium," or "serious." For this research, the LHS's categorization of violation types is used to reflect the nature of responsive enforcement, though these categorical labels may be misleading. For example, "unintentional" violators may not necessarily be less severe offenders in the real world; one could commit a high severity offense unintentionally. "Minor" violations could be intentionally committed. Yet, this research uses the categorization of the LHS that define "unintentional" violators as inspectees whose offenses are the least severe. The model assumes the verbiage of the LHS despite its shortcomings.

4. Absorbance capacity indicator

This indicator is a logit-normal distribution that characterizes the capacity of the inspectee to absorb the enforcement intervention, which means that they are able to sustain their business even after the intervention. This variable, *absorbance-capacity*, has a value between zero and one. An *absorbance-capacity* indicator close to one indicates that the inspectee has a high capacity to sustain their business even after the enforcement intervention, while an *absorbance-capacity* indicator closer to zero means that the inspectee likely cannot bear the consequences of the intervention. The logit-normal distribution was chosen as it always outputs a value between zero and one and allows the user to define the shape of the distribution by varying the mean (μ) and standard deviation (σ). The logit-normal probability density function and the variable's logit are given by Equations 3.1a and 3.1b, respectively. Note that this indicator is not a phenomenological parameter as there is no empirical data on how inspectees respond to enforcement.

$$f x(x;\mu,\sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{\frac{(\mathrm{logit(x)}-\mu)^2}{2\sigma^2}} \frac{1}{x(1-x)} \tag{3.1a}$$

$$\mathrm{logit}(x) = log\left(\frac{1}{x(1-x)}\right) \tag{3.1b}$$

Figure 3.1 shows various logit-normal probability density functions given different values of μ and σ.

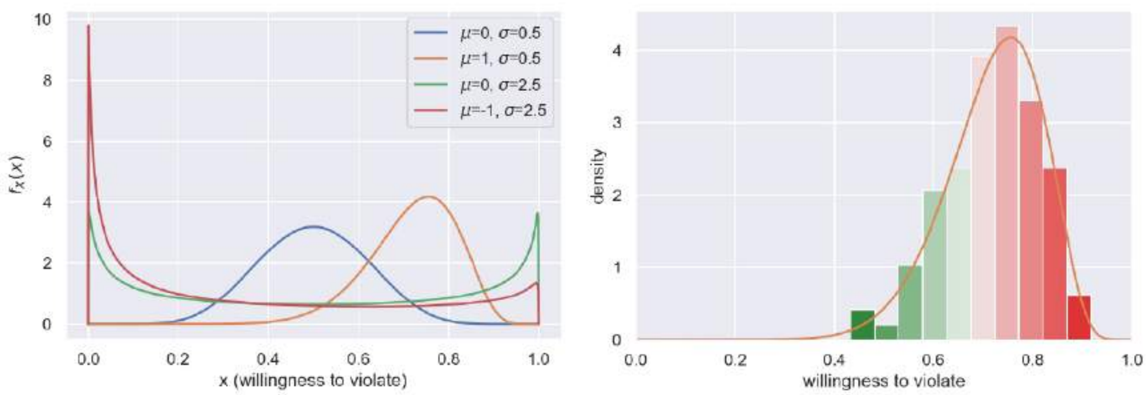

(a) Examples of logit-normal distributions.

(b) Example of a sample with $\mu = 1$ and $\sigma = 0.5$.

Figure 3.1: Logit-normal distributions
with different values of μ and σ (Knol, 2021).

3.3 PABMI: Time Scale

The PABMI is a time simulation where inspectee agents decide on their behavior in each time step called *tick*. The PABMI has a user-definable parameter called the *length-of-run* where the model will run for the defined number of ticks. Each tick represents one cycle of inspections. At each tick, inspections are always conducted. According to the Inspectieview Binnenvaart data, inspections can occur multiple times a year or as infrequent as once in two years, depending on the agency. The PABMI does not assign a specific time period to each tick, but rather uses each tick to represent one round of inspections.

3.4 PABMI: Model Process Overview

The PABMI follows the simplified model narrative as shown in Table 3.3.

1. Setup Procedure
1a) Create inspectee agents (see Section 3.2)
2. Go Procedure
2a) Inspectors choose inspection candidates (Section 3.4.1)
2b) Inspectors inspect candidates and perform enforcement (Section 3.4.2)
2c) Inspectees update their behavior (Section 3.4.3)

Table 3.3: Simplified model narrative for the PABMI.

The setup procedure is detailed in Section 3.2 when the inspectee agents are created. The rest of this section outlines the Go procedures which happens at every time step and consists of three main sub-models: Sub-model 2a: Inspectors choose inspection candidates, Sub-model 2b: Inspectors inspect candidates and perform enforcement, and Submodel 2c: Inspectees update their behavior.

3.4.1 Go Procedure: Inspectors choose inspection candidates

In this sub-model, the inspectee agents are chosen for inspection. This sub-model refers to Go Procedure #2a in Table 3.3. There are five strategies that inspectors can use to select inspection candidates: 1. All random (AR), 2. Risk-based using the non-compliant record (RbNC), 3. Risk-based using offense severity (RbOS), 4. Mix of random and risk-based using non-compliant record (MRRbNC), and 5. Mix of random and risk-based using offense severity (MRRbOS) (see Table 3.4).

#	Inspection Strategy	How Inspectees are Chosen for Inspection
1	All Random (AR)	Select inspectees randomly
2	Risk-based: Highest non-compliant record (RbNC)	Select inspectees with the highest track record of past non-compliant inspections
3	Risk-based: Highest offense severity (RbOS)	Select inspectees with the highest offense severity (criminal offenders)
4	Mix of Random and Risk-based: based on non-compliant record (MRRbNC)	Select a proportion of inspectees for random inspection (based on percent-random-inspections, user-definable parameter); for the remaining number of inspections, select inspectees with the highest track record of past non-compliant inspections
5	Mix of Random and Risk-based: based on offense severity (MRRbOS)	Select a proportion of inspectees for random inspection (based on percent-random-inspections, user-definable parameter); for the remaining number of inspections, select inspectees with the highest offense severity (criminal offenders)

Table 3.4: PABMI inspection strategies.

The capacity of the inspectorate is determined by the number of inspectors (*n-inspectors*) multiplied by the number of inspections one inspector can conduct in one inspection cycle (*n-inspections-per-inspector*). This indicates the total number of inspections that are conducted in one time step. These two parameters are user-definable and can be adjusted with a slider on the Netlogo interface.

3.4.2 Go Procedure: Inspectors inspect candidates and perform enforcement

In this sub-model, the inspection candidates chosen in Sub-model 2a (Section 3.4.1) are inspected, and enforcement interventions are performed. This sub-model refers to Go Procedure #2b in Table 3.3. For inspectees found to be compliant, their state variables are updated as follows:

- Number of compliant inspections (*n-compliant*) increases by 1
- Number of times inspected (*n-times-inspected*) increases by 1
- Compliance rate (*compliance-rate%*) is updated by dividing the *n-compliant* by the *n-times-inspected*

For inspectees found to be non-compliant, their state variables are updated as follows:

- Number of non-compliant inspections (*n-noncompliant*) increases by 1
- Number of times inspected (*n-times-inspected*) increases by 1
- Compliance rate (*compliance-rate%*) is updated by dividing the *n-compliant* by the *n-times-inspected*
- Compliance status (*compliance*) is updated according to the type of offense ("unintentional," "conscious," or "criminal")
- Severity of offense (*offense-severity*) is updated according to compliance status (see Table 3.2).

The user can define the enforcement strategy as one of two options: 1) Standard enforcement (SE), or 2) Responsive enforcement (RE). SE represents the pre-LHS inspection strategy where all non-compliant inspectees received enforcement interventions of a standard and similar level of severity (see Secion 1.2.1.1). With SE, inspectors largely gave low-level enforcement interventions like warnings and administrative action (see Table 2.8 and Section 2.1.2). On the contrary, RE refers to the responsive regulatory strategy where the severity of enforcement given to non-compliant inspectees is commensurate with the severity of the offense (see Section

1.2.1.1). This definition of RE is based on the LHS which dictates that the severity of enforcement should be determined based on the offense severity. While this is admittedly a narrow interpretation of RE, the model only depicts RE based on the offense severity to be consistent with the LHS (see Figure 3.2). Yet, in responsive regulation, enforcement doled out to non-compliant inspectees can be determined based on a number of different factors (see Section 1.2.1.1).

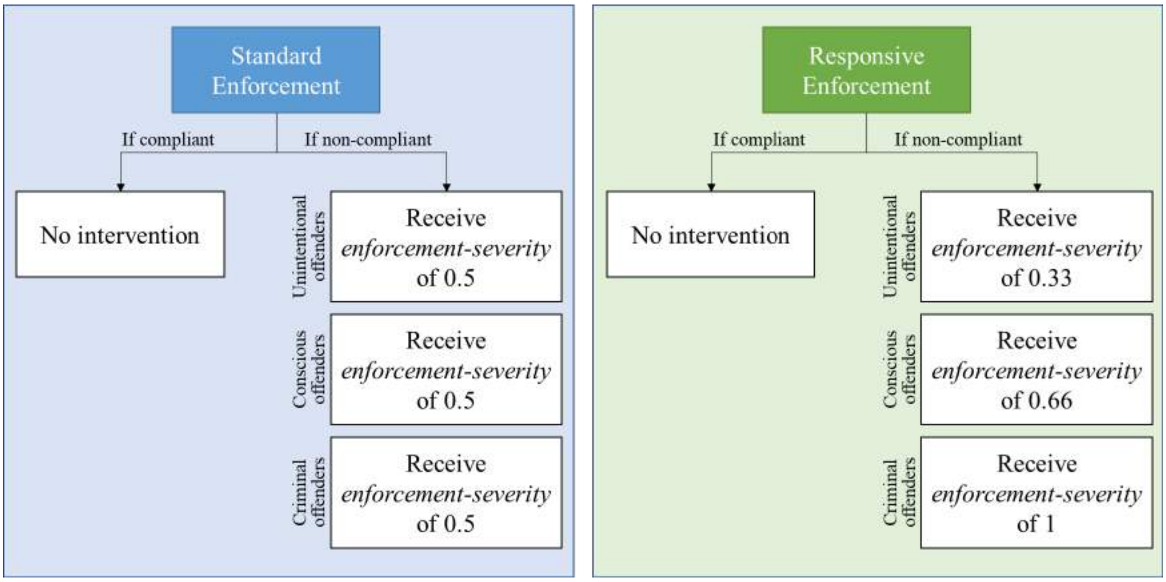

Figure 3.2: PABMI standard versus responsive enforcement strategies.

The *enforcement-severity* state variable is an indicator between zero and one that reflects the severity of the enforcement. An *enforcement-severity* value close to one (zero) means that the severity of enforcement is high (low). For example, unintentional offenders commit smaller-scale and low-impact offenses (see Section 3.4.1). With RE, their enforcement intervention should be the least severe like a warning (see Figure 1.2); therefore, the *enforcement-severity* indicator will be closest to zero (see Figure 3.2). Conscious violators commit offenses with a medium impact; the severity of the enforcement intervention should be moderate like an administrative fine and temporary shutdown (see Figure 1.2). The criminal offenders should be given the most severe enforcement intervention such as a criminal penalty, police report, cease and desist order, or permanent shutdown (see Figure 1.2).

As previously mentioned, the severity of enforcement under SE does not change with compliance level. Therefore, the *enforcement-severity* indicator is set to a standard, neutral value for all types of non-compliant inspectee (see Figure 3.2).

3.4.3 Go Procedure: Inspectors update their behavior

This sub-model refers to Go Procedure #2c in Table 3.3. The behaviors of the inspectees are determined based on three categories of observed behavioral phenomena: 1) Peer pressure, 2) Inspections, and 3) Enforcement (see Table 2.3). These reactions were observed by inspectors and discovered through analysis of the Inspectieview Binnenvaart data (see Section 2.1.2). These reaction types can be switched on and off by the user on the Netlogo interface to isolate each behavior type in the PABMI simulation and to test different combinations of reactions.

Reaction to peer pressure

When inspectees react to peer pressure, they mimic the most occurring offense severity level among their neighbors that are within a particular user-definable *radius*. In other words, they check for the mode of the *offense-severity* of other inspectees within the *radius* and adopt that behavior. This behavior is not only observed by ILT's inspectors (see Section 2.1.2) but also based on the behavioral theory that individual decisions of actors in a network are driven by pressures from other actors (Durlauf, 1999; Keizer et al., 2008; also see Section 1.2.2.2). When inspectees see their peers violating the law, they too might be more likely to violate. Because there is no hard empirical data on the effects of peer pressure, this micro-interaction had to be conceptualized from theory and qualitative data. While this approach reflects a more traditional approach to ABM rather than a phenomenological approach (see Table 1.4), this behavior was included because inspectors observe peer pressure effects in the inspection environment. The PABMI applies behavioral theories to fill in the gaps in empirical data.

Reaction to inspections

When inspectees react to inspections, some of them choose to escalate their behavior (become higher level offenders), deescalate their behavior (become lower level offenders), or stay the same (see Table 2.3). This phenomena can be observed in the Inspectieview Binnenvaart data and is included into the model as user-definable parameters called *percent-escalate* and *percent-deescalate*. Escalatory inspectees raise their compliance status to a higher level of severity when they are inspected. Conversely, deescalatory inspectees lower their compliance status to a lower level of severity when they are inspected (see Table 3.5). These variables use the statistical patterns found in empirical data (see Section 2.1.2) to conceptualize this behavior in the ABM.

		Escalatory Inspectees		Deescalatory Inspectees	
Current Compliance Level		**New Compliance Level**	**New Offense Severity**	**New Compliance Level**	**New Offense Severity**
Unintentional	→	Conscious	20	Compliant	0
Conscious	→	Criminal	30	Unintentional	10
Criminal	→	Criminal	30	Conscious	20

Table 3.5: PABMI escalatory versus deescalatory behavior.

Reaction to enforcement

When inspectees react to enforcement, they assess whether or not they can sustain the severity of the enforcement before choosing their next behavior. If the *absorbance-capacity* of the inspectee is greater than or equal to the *enforcement-severity*, they will continue to violate at the same offense severity level. If their *absorbance-capacity* is less than the *enforcement-severity*, it means they cannot afford to continue violating, so they will decide to comply. This simulates the learning effect of the inspectees; if they learn that they can bear the cost of the enforcement intervention, they continue violating. The model assumes that only inspected agents can react to enforcement, as only non-compliant inspectees are given enforcement interventions.

Similar to the reaction to peer pressure, the reaction to enforcement is based on behavioral theories. Specifically, SE is based on the pre-LHS "Interventieladder" policy, and RE is based on the responsive regulatory theory of the LHS (see Section 1.2.1.1). While this is not a purely phenomenological approach, this behavior was included in the PABMI, as the effectiveness of the LHS for increasing compliance has not yet been tested. Investigating the effects of the LHS on macro-level system behavior allows the ILT to explore the conditions under which a reaction to enforcement produces better outcomes.

To summarize, Figure 3.3 shows the flow of the PABMI.

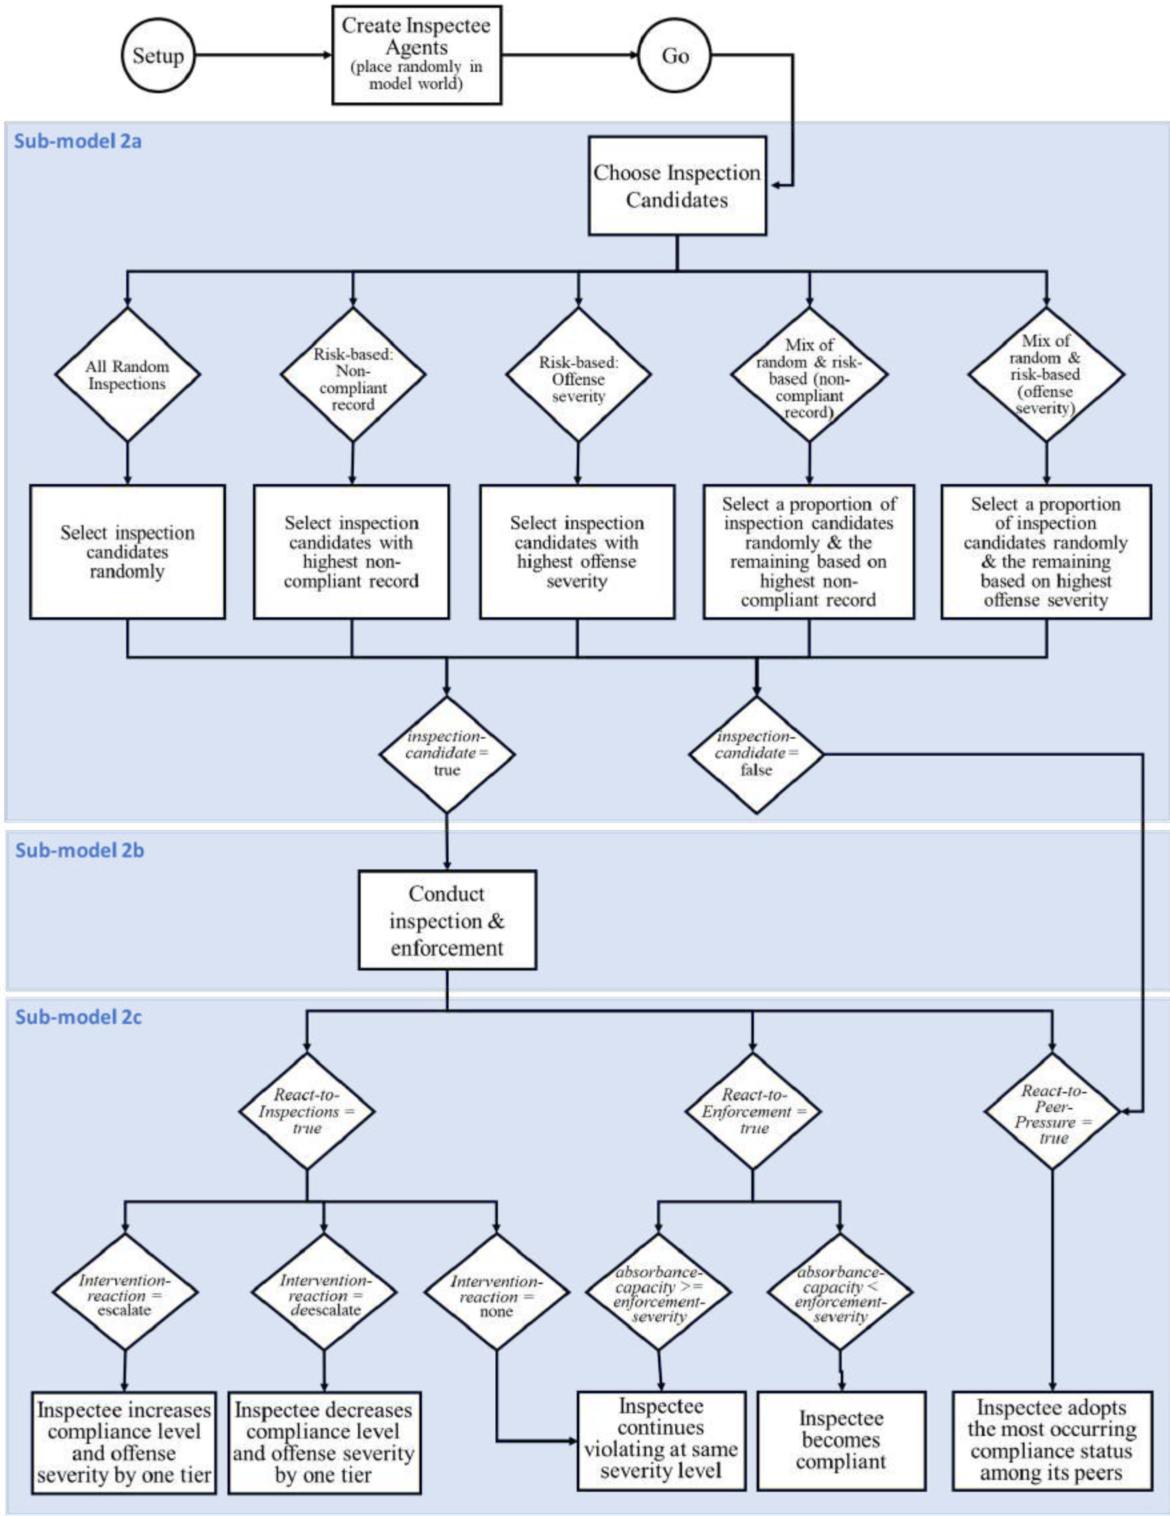

Figure 3.3: PABMI model flow diagram.
The circles indicate the start of an event while the rectangles denote the processes in the model. The diamonds determine the conditions that need to be fulfilled to continue down the pathway indicated by the arrows.

Chapter 4

Implementation of the Phenomenological ABM for Inspections (PABMI)

4.1 PABMI: Assumptions

The PABMI assumes the following:

- Inspection candidates are only inspected once in each time step. There are no situations where the inspection candidates are inspected more than once in one inspection cycle. In reality, inspectees are inspected at various rates. For example, some get inspected every 32 weeks while others are inspected every 109 weeks (Meester, 2021).
- All inspectees are available for inspection at every time step. This may not be realistic, as some inspectees may be in motion and may not be available for inspection in a given inspection cycle.
- Inspectees have a binary choice of either complying or not complying. Also, there is no option for the inspectee to violate multiple times in one time step.
- All non-compliant inspectees are caught. This may not be realistic, as non-compliant inspectees in the real world could thwart inspections so that they are found to be compliant when inspected.
- All inspected non-compliant inspectees are given enforcement interventions. There is no situation where an inspectee is inspected, deemed to be non-compliant, and not given an enforcement intervention.
- Deescalatory and escalatory inspectees decrease and increase their offense severity by one compliance level in each time step, respectively. In other words, if an inspection candidate reacts to the inspection, they change their behavior by only one compliance level in each time step. The only circumstance in which they will skip compliance levels is if their absorbance capacity does not allow them to sustain their level of non-compliance.
- All inspection candidates are inspected at a uniform level of thoroughness. In reality, each inspection agency conducts inspections at different levels of thoroughness. However, there is little transparency into the procedures they

use to conduct the inspections. Therefore, by assuming that all inspection candidates are inspected in the same way, the PABMI simulates a spreading of inspection risk across the inspection agencies.

- The PABMI does not consider specific regulations that are violated.
- The PABMI assumes that minor, medium, and serious violations as indicated in the Inspectieview Binnenvaart dataset are unintentional, conscious, and criminal offenses as defined by the LHS, respectively.

4.2 PABMI: Model Calibration

The PABMI is calibrated to the Inspectieview Binnenvaart data as described in Section 2.1.2, which is crucial for a phenomenological approach to ABM (see Table 1.4). Table 4.1 shows the default parameter values for the PABMI.

Parameter	**Value**	**Source**
Number of inspectees	7878	Table 2.4
Length of run	1000	Assumption
% of inspectees who are compliant	39.6%	Figure 2.1
% of the non-compliant population who are unintentional offenders	20.8%	Table 2.6b
% of the non-compliant population who are conscious offenders	76.4%	Table 2.6b
% of the non-compliant population who are criminal offenders	2.8%	Table 2.6b
Average number of violations for unintentional offenders	3.99	Table 2.7
Average number of violations for conscious offenders	14.67	Table 2.7
Average number of violations for criminal offenders	0.54	Table 2.7
Number of inspectors	100[a]	Table 2.4
Number of inspections per inspector	13.13[a]	Table 2.4
% of random inspections	20%[b,c]	Knol, 2021
% of the non-compliant population who are escalatory	2.97%	Figure 2.2
% of the non-compliant population who are deescalatory	6.98%	Figure 2.2
Logit normal mu (μ) parameter for absorbance capacity	1[c]	Assumption
Logit normal sigma (σ) parameter for absorbance capacity	0.5[c]	Assumption
Radius of peers	10[c]	Assumption

[a] The capacity of the inspectorate is the number of inspectors multiplied by the number of inspections. Because the ILT historically conducted 1,313 inspections (see Table 2.4), these parameters were adjusted to total 1,313 inspections when multiplied together.

[b] Knol (2021) found 20% to be the optimal random share of inspections; however, this parameter in her model reflects the learning effects of the inspectors over time. The PABMI assumes that the inspectors have knowledge of the track record of inspectees, as this was determined based on statistical patterns found in empirical data. 20% is used as a starting point for the value of this parameter, but is later adjusted based on the sensitivity analysis.

[c] These parameters are not determined phenomenologically, but rather rely on theories on the target population to calibrate the interactions between agents. This reflects a more traditional approach rather than a phenomenological approach to ABM (see Table 1.4). The reasoning behind the inclusion of these parameters is discussed in Section 3.4.

Table 4.1: PABMI default parameter values.

4.3 PABMI: Reporting Outcomes of Interest

The PABMI model reports the average compliance rate of the whole inspectee population and the breakdown of the inspectee population by compliance level. The average compliance rate represents the average ratio of compliance to the number of times an agent is inspected. The breakdown of the inspectee population includes the percentage of compliant inspectees, unintentional violators, conscious violators, and criminal violators.

4.4 PABMI: Uncertainties

The PABMI has the following uncertainties:

- The Inspectieview Binnenvaart data quality is low. The violations and enforcement interventions have different systems of classification and no standardized definition. There are also unclassified records of violation and enforcement. The model disregards these unclassified records and takes the population breakdown of the violation types without counting the unclassified records.
- Inspection agencies have different capacities for enforcement, and there is no standardized enforcement strategy across agencies. For example, only the police can file a police report. The enforcement intervention is largely up to the discretion of the inspector, which makes it tough to accurately predict the type of enforcement that will be executed.
- The effectiveness and validity of the LHS has yet to be proven. Yet, the PABMI conceptualizes the mechanism of RE based on the LHS to investigate conditions under which the LHS is effective (see Section 3.4). This must be validated in the real world.
- The absorbance capacity parameter is an indicator that was created in an attempt to quantify inspectees' tolerance of enforcement interventions. There is no existing data on how inspectees react to enforcement. Therefore, this indicator was created to characterize the distribution of inspectees at varying capacities to withstand enforcement.
- There are a plethora of other external factors that could impact compliance behavior, and it is impossible to know or model all of them. The PABMI is restricted to examine only three influencing forces of behavior: peer pressure, inspections, and enforcement (see Table 2.3 and Section 3.4).

Chapter 5

Results of the Phenomenological ABM for Inspections (PABMI)

5.1 PABMI: Sensitivity Analysis and Structural Validation

Using NetLogo's BehaviorSpace and Python 3.8, a global sensitivity analysis was conducted on the PABMI to assess the sensitivity of input parameters in a systematic way. This analysis shows the degree of influence a model parameter has on the model output. The aim of the sensitivity analysis is to learn about the model and test the boundaries of its usefulness. Sensitivity analyses help to "[identify]...important thresholds," thereby "[disciplining] the dialogue about options and make unavoidable judgments more considered" (Epstein, 2008). Since the PABMI combines empirical and theoretical parameters, the sensitivity analysis illuminates the intricacies of the model dynamics.

Five parameters were varied (see Table 5.1) while keeping default values constant for other parameters (see Table 4.1). Each of these varied parameters was simulated 50 times over 1000 time steps. With the exception of the sensitivity analysis on the percent of random inspections, the sensitivity analyses were conducted using an all random (AR) inspection strategy.

Parameter	Min Value	Max Value	Varied by
Initial share of compliant inspectees ($\%_{init,comp}$)	10%	70%	10%
% of random inspections ($\%_{rand}$)	10	90	10
Logit mu (μ)	0.1	1.0	0.1
Logit sigma (σ)	0.1	1.0	0.1
Radius of peers (r_{peers})	5	50	5

Table 5.1: Varied parameters for the PABMI sensitivity analysis.

Structural validation shows the boundaries of model parameters where the model no longer outputs reasonable results. According to ILT experts, a non-compliant proportion greater than 50% of the population is a realistic threshold for any sector. A non-compliant share greater than 50% indicates that regulations may be too expensive to comply with or they are misaligned with real-life circumstances. In this case, regulations are usually edited so that more inspectees can meet regulatory standards; enforcement is not pursued. Therefore, a threshold of 50% of compliant inspectees is used to structurally validate the model parameters.

Structural validation was the chosen method of validation for this research given time and resource constraints of the research and limited access to inspectors. More importantly, other methods such as validation with empirical data or cross validation with Knol's ABM (see Section 2.1.1) are not viable for this research. Since the empirical data was used to create the PABMI, it cannot be used for validation. In future research, larger longitudinal datasets of the sector can be divided into train and test sets so that a portion can be used to build the model while the other is used to validate the model. Knol's ABM has significantly different conceptualization, assumptions, and purpose; therefore, a cross validation with that model is not reasonable (see Section 2.1.1). Future research should conduct expert validation where inspectors are consulted on the validity and usability of the PABMI.

5.1.1 Varying Initial Share of Compliant Inspectees ($\%_{init,comp}$)

The results of the global sensitivity analysis show that the outcomes of interest (see Section 4.3) are sensitive to changes in the initial percentage of compliant inspectees ($\%_{init,comp}$), particularly between 30% and 50% (Figures 5.1 and 5.2). Across all values of $\%_{init,comp}$, RE consistently yields a higher average compliance rate among inspectees when all inspection candidates are identified randomly (AR inspection strategy; see Table 3.4).

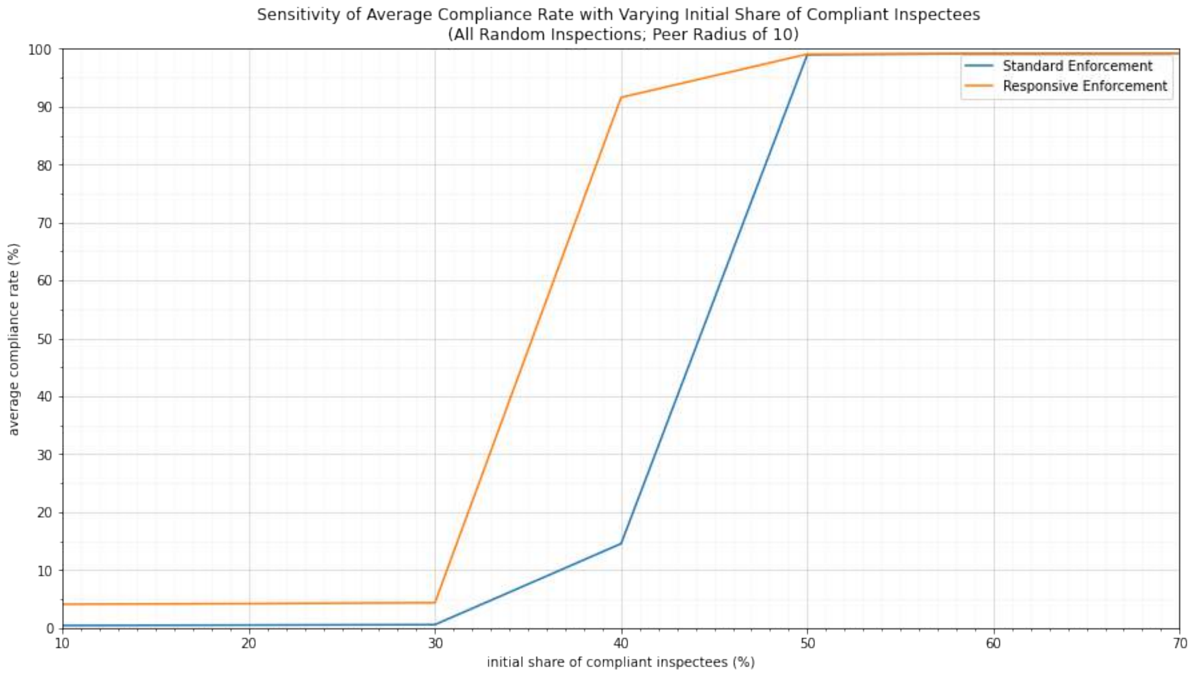

Figure 5.1: Sensitivity of the average compliance rate with varying $\%_{init,comp}$.

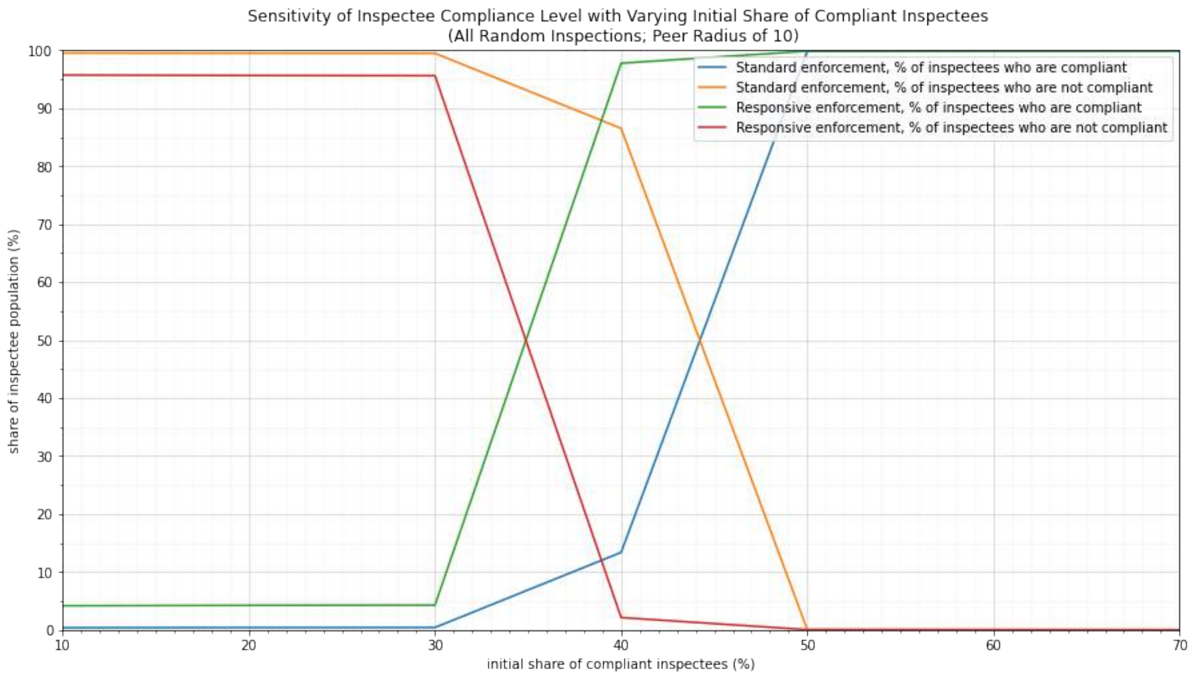

Figure 5.2: Sensitivity of the compliance status breakdown with varying $\%_{init,comp}$.

With both SE and RE, there is an overwhelming larger number of conscious offenders on average compared to the unintentional or criminal offenders (Figures 5.3a and 5.3b). The inland shipping data shows that conscious violators make up the largest proportion of non-compliant inspectees (Table 2.6b); therefore, the likelihood that an inspectee has a neighbor who is a conscious violator is high. With peer pressure effects, this means that there is a high likelihood that an inspectee is influenced by a conscious violator. In addition, since the severity of SE is constant, inspectees with an absorbance capacity greater than 0.5 can consistently continue violating at their current compliance level (see Figure 3.2 for the model conceptualization of inspectees' reaction to SE versus RE). If many are influenced by conscious violators, the population of non-compliant inspectees will become mainly conscious violators. The effect of peer pressure coupled with the ability of inspectees to cope with the severity of SE produces a population of inspectees where the behavior of conscious violators are most often adopted.

SE requires a higher $\%_{init,comp}$ (at least 40%) compared to the RE (at least 30%) to maintain a majority compliant population (Figure 5.3). This indicates that SE relies more on positive peer pressure to maintain compliance than RE, where the variation of the enforcement severity influences inspectees to become compliant more so than peer pressure.

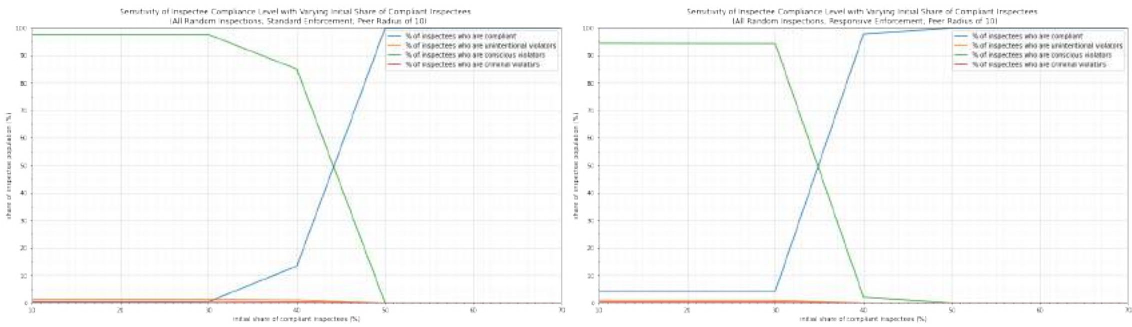

(a) Standard enforcement. (b) Responsive enforcement.

Figure 5.3: Sensitivity of the compliance level breakdown with varying $\%_{init,comp}$.

Structural Validation

With SE, the $\%_{init,comp}$ should be at least 30% and below 50% for the model to be valid (Figure 5.3a). Below 30%, there are no compliant inspectees. Above 50%, all inspectees are compliant. Neither of these situations are structurally valid. With RE, the $\%_{init,comp}$ should be between 30% and 40% (Figure 5.3b). Therefore, further sensitivity analyses were conducted where $\%_{init,comp}$ was varied in smaller increments (2.5%) between 30% and 50% to show its reasonable limits more precisely (Figure 5.4 and 5.5).

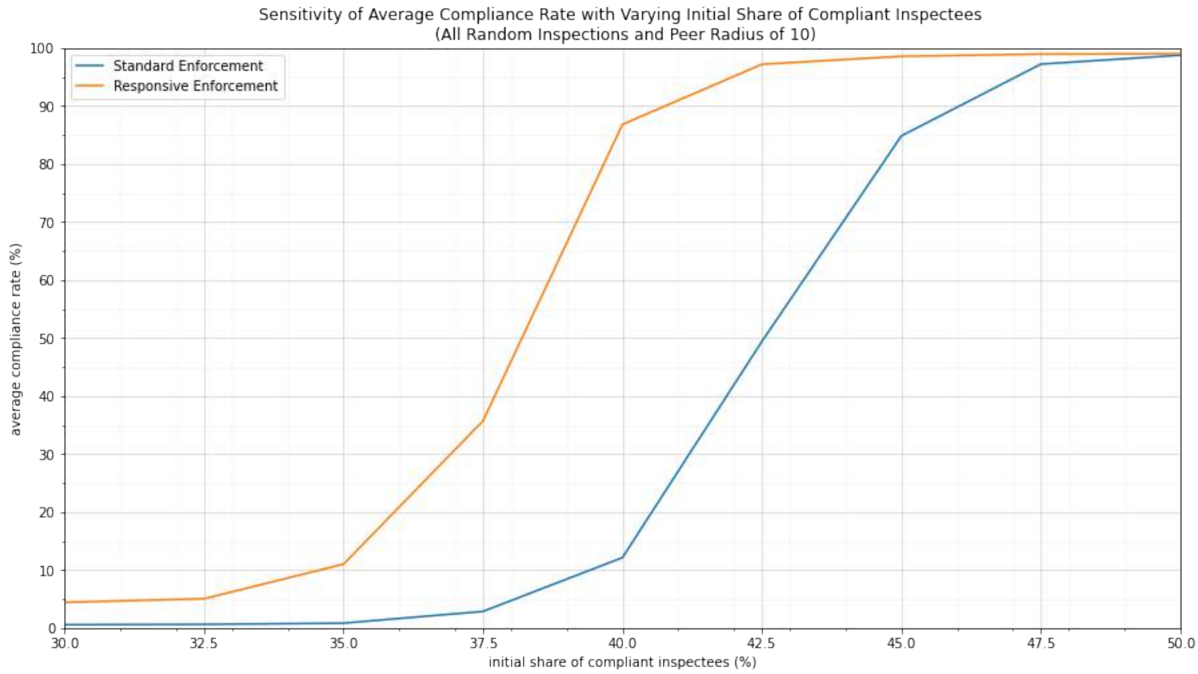

Figure 5.4: Sensitivity of the average compliance rate ($\%_{init,comp}$: 30-50%).

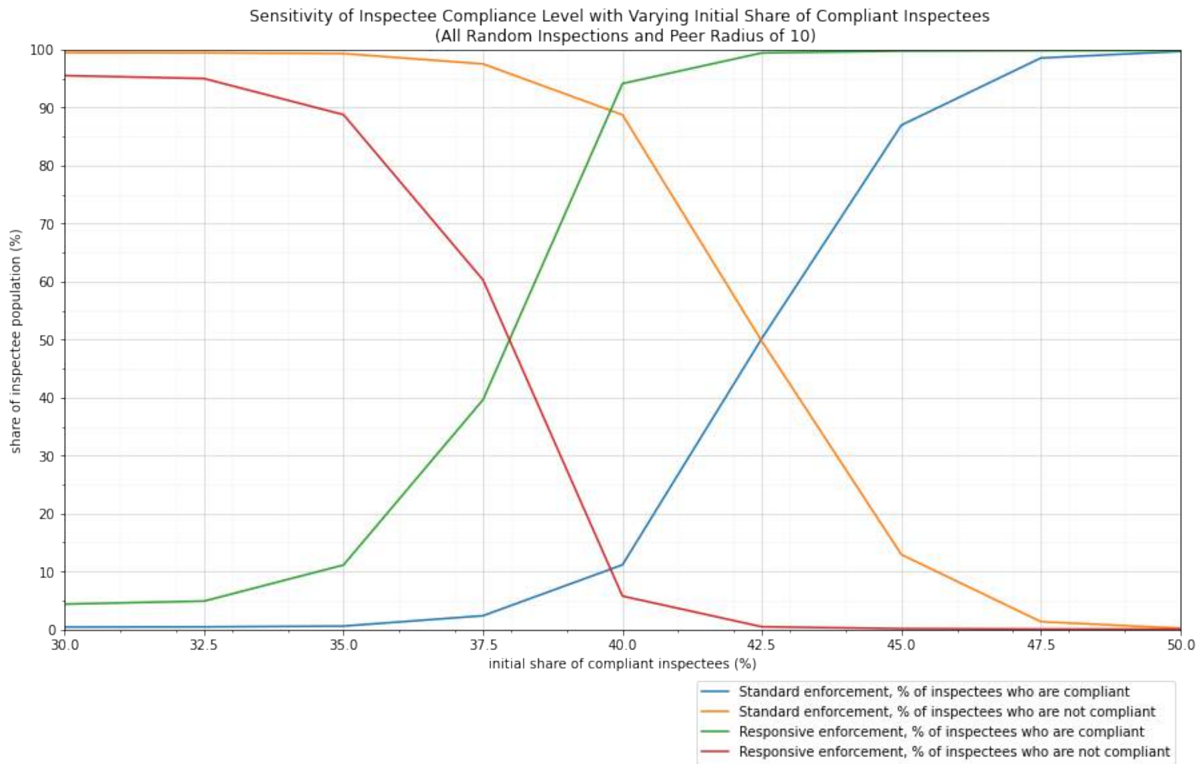

Figure 5.5: Sensitivity of the compliance status breakdown ($\%_{init,comp}$: 30-50%).

Standard Enforcement

With SE, an initial population of 42.5% compliant inspectees yields an equal proportion of compliant and non-compliant inspectees in the PABMI (Figure 5.6).

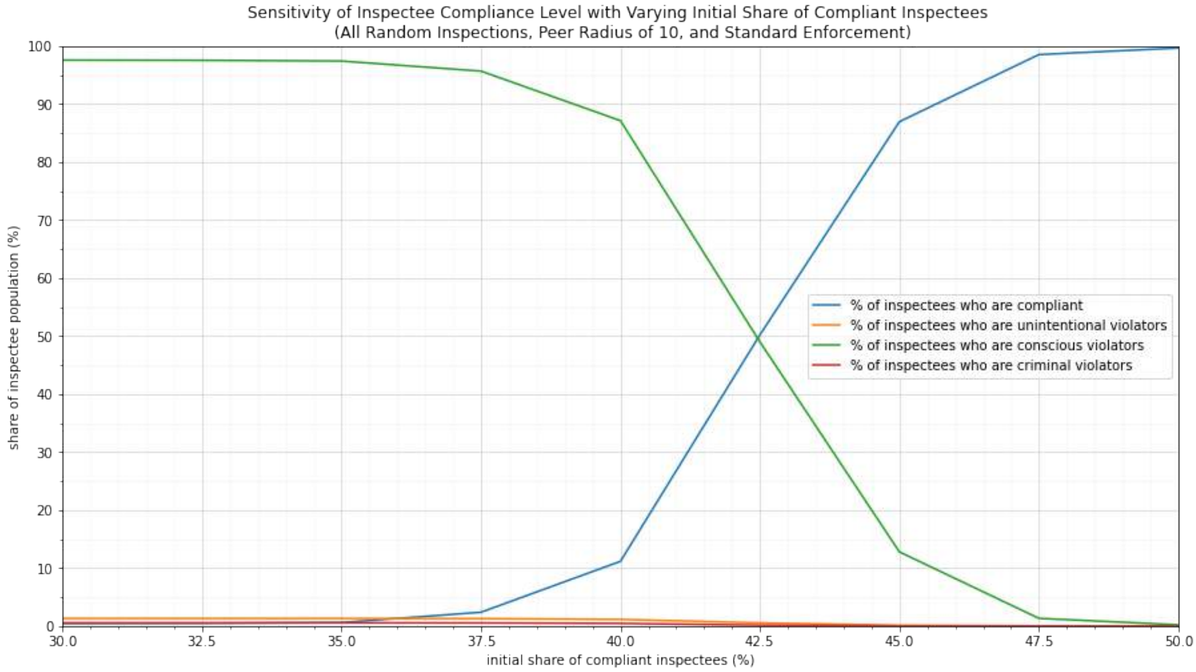

Figure 5.6: Sensitivity of the compliance level breakdown ($\%_{init,comp}$: 30-50%; SE).

The reasonable range of the $\%_{init,comp}$ is between 42.5% and 47.5% (Figure 5.7; also see Figure A.1 for the average compliance rate over time). Outside of this range, the compliance rate of the inspectees converges unrealistically towards zero or 100, respectively, or produces too many non-compliant inspectees to be structurally valid.

In the inland shipping sector, the initial share of compliant inspectees is 39.6% (see Table 4.1). The results of the PABMI sensitivity analysis show that this percentage is too low to create a population of majority compliant inspectees with SE. This means that given a peer radius of 10 and an AR inspection strategy, SE as defined in the model conceptualization (see Section 3.4) is insufficient for increasing compliance rates given the initialization of the model parameters based on empirical data.

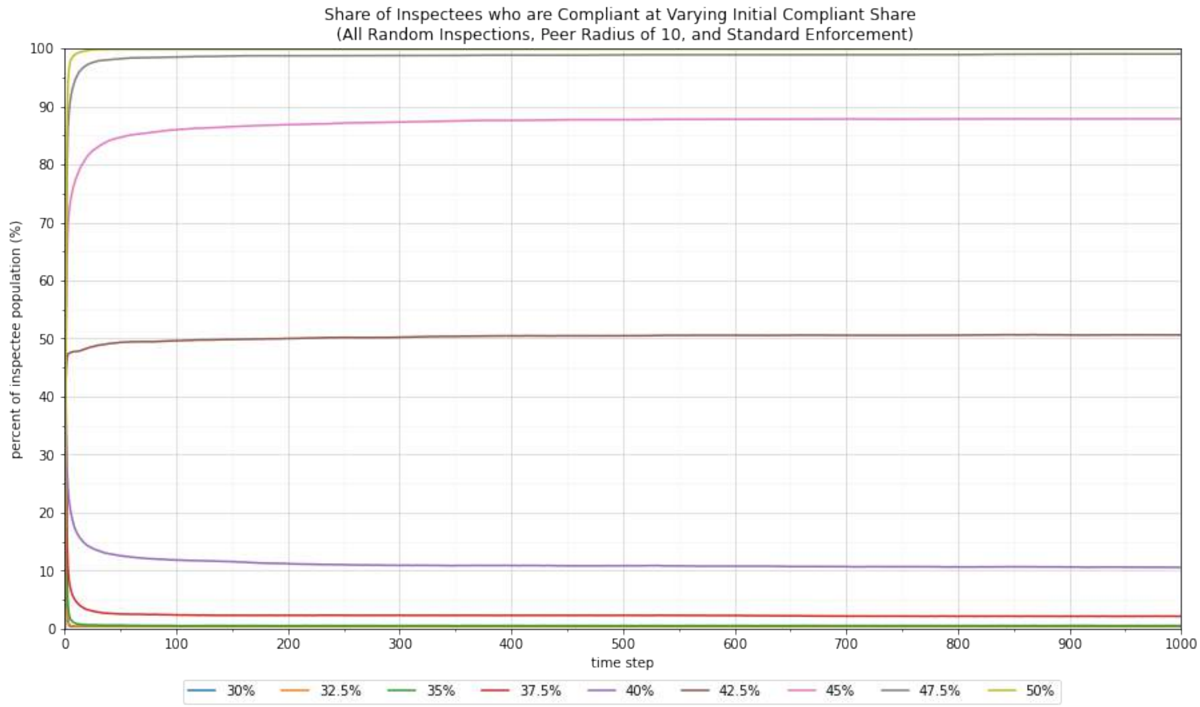

Figure 5.7: Share of compliant inspectees over time with varying $\%_{init,comp}$ (SE). For a breakdown by compliance level, see Figures A.2 and A.3.

Responsive Enforcement

A smaller share of compliant inspectees is required to reach an equilibrium between the population of compliant to non-compliant inspectees with RE compared to SE (Figures 5.6 and 5.8). A minimum initial population of 38% of compliant inspectees is required for RE compared to 42.5% for SE (Figure 5.8). This model behavior is a result of the inability of more non-compliant inspectees to absorb the severity of RE compared to SE, especially at higher offense severity levels. RE requires less reliance on positive peer pressure to yield a majority compliant population compared to SE. It is worth noting that both the effects of peer pressure and enforcement are based on inspector's anecdotal evidence and behavioral theories; they were not derived from empirical data (see Section 3.4.3). Therefore, the conclusion that RE provides more leeway for non-compliant peer pressure still needs to be verified in the real world. In sum, the structurally valid range of the $\%_{init,comp}$ is between 38% and 40% (Figure 5.9; see Figures A.4 and A.5 for the average compliance rate and share of non-compliant inspectees, respectively).

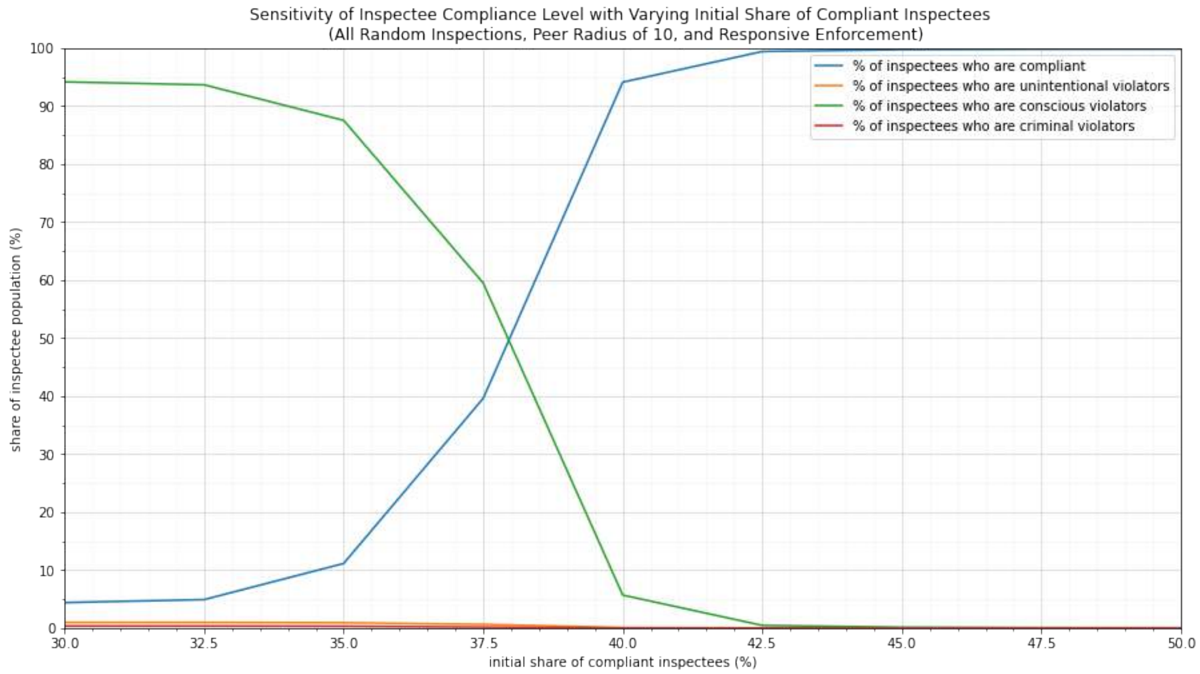

Figure 5.8: Sensitivity of the compliance level breakdown ($\%_{init,comp}$: 30-50%; RE).

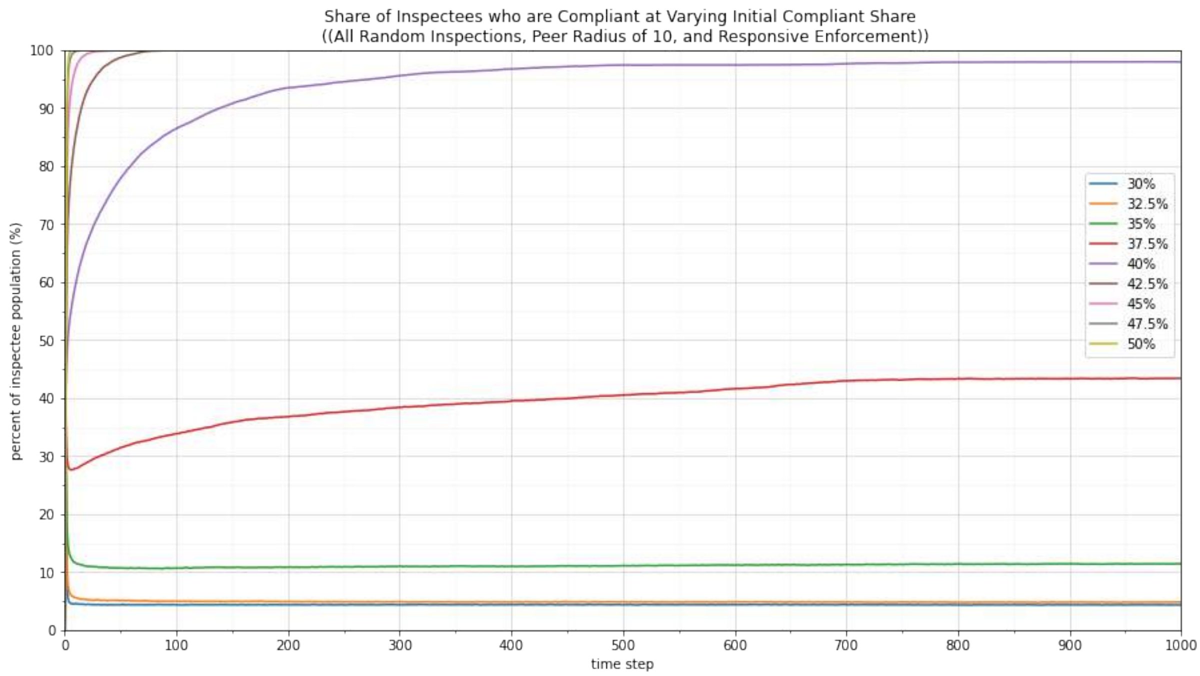

Figure 5.9: Share of compliant inspectees over time with varying $\%_{init,comp}$ (RE). For a breakdown by compliance level, see Figures A.5 and A.6.

Path Dependency

Within valid ranges of the $\%_{init,comp}$, the averaged values of the outcomes of interest vary for both SE and RE (Figure 5.10, see Figure A.7 for the spread of the share of compliant inspectees). With SE, the outcomes of interest show larger variation and stronger path dependency when the $\%_{init,comp}$ equals 42.5% compared to when it equals 45% or 47.5% (Figure 5.11). This shows that with less compliant inspectees to exert positive peer pressure, the placement of the inspectees in relation to one another can produce peer pressure effects that alter the behavioral pathway of the whole population. A higher compliant share relies less on the placement of inspectees for the population to become more compliant over time.

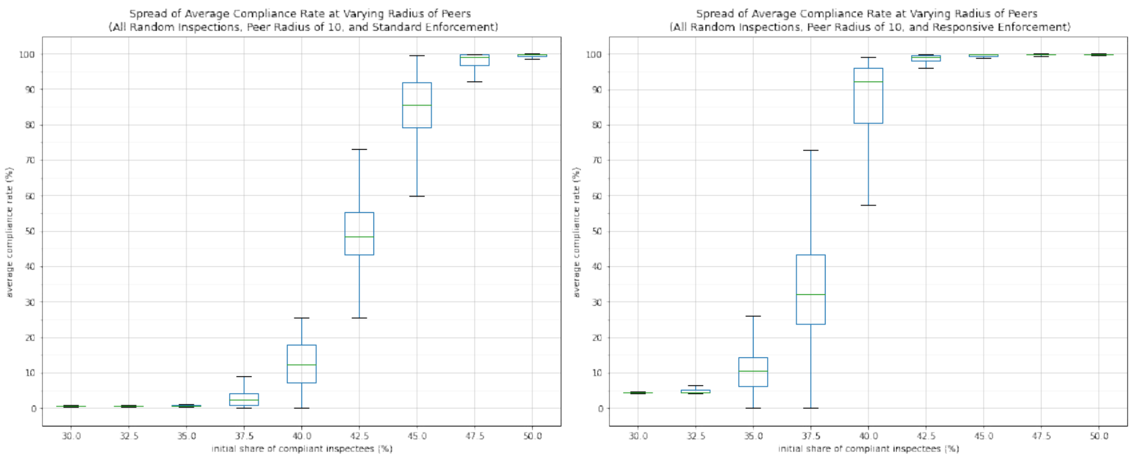

(a) Standard enforcement.

(b) Responsive enforcement.

Figure 5.10: Spread of the average compliance rate with varying $\%_{init,comp}$.

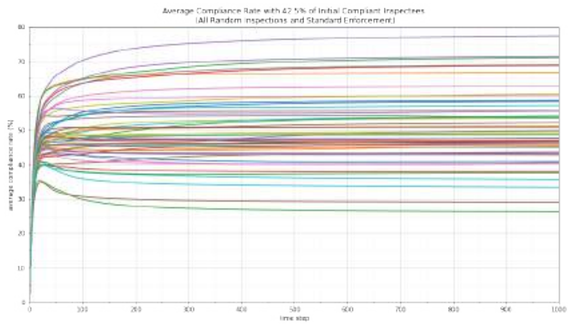

(a) 42.5% initial compliant inspectees.
For a breakdown by compliance level, see Figure A.8.

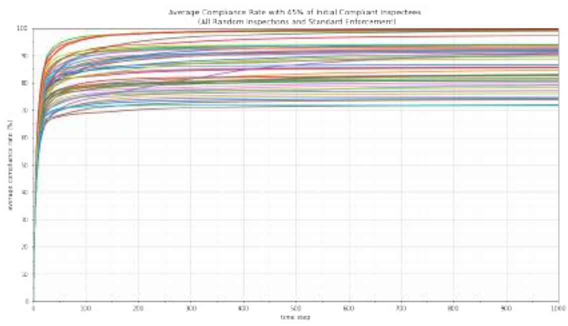

(b) 45% initial compliant inspectees.
For a breakdown by compliance level, see Figure A.9.

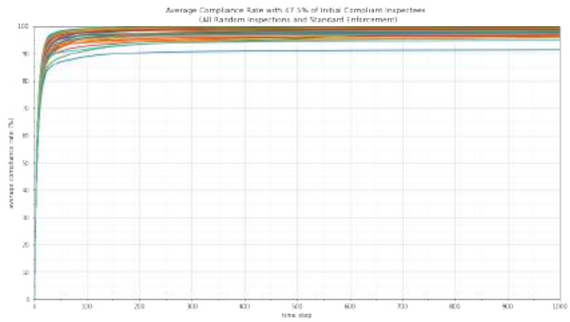

(c) 47.5% initial compliant inspectees.
For a breakdown by compliance level, see Figure A.10.

Figure 5.11: Average compliance rate results ($\%_{init,comp}$: 42.5%, 45%, 47.5%; SE).

With RE, a similar pattern can be found. There is more variation in the simulation pathway for the lower $\%_{init,comp}$ value in the valid range than the higher one (Figure 5.12).

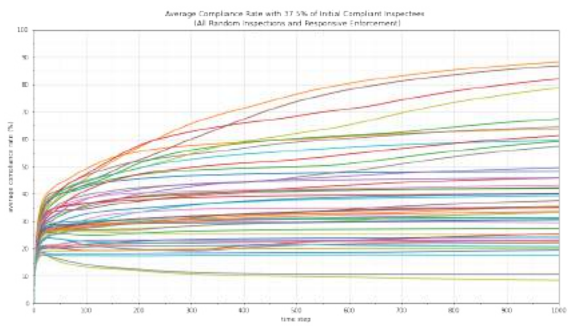

(a) 37.5% initial compliant inspectees.
For a breakdown by compliance level, see Figure A.11.

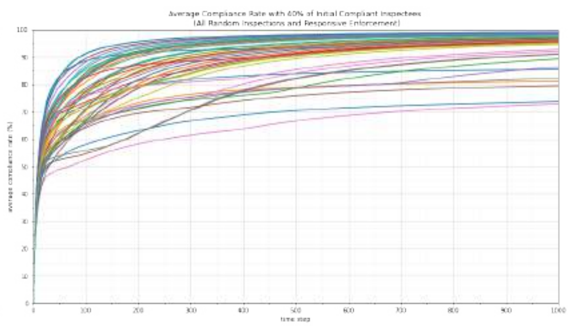

(b) 40% initial compliant inspectees.
For a breakdown by compliance level, see Figure A.12.

Figure 5.12: Average compliance rate results ($\%_{init,comp}$: 37.5% & 40%; RE).

The outcomes of interest have stronger path dependency and larger variation with RE than SE within valid ranges of the $\%_{init,comp}$. With RE, more severe non-compliant actors cannot absorb the severity of enforcement when they are inspected, so they are more likely to become compliant over time. This portion of non-compliant inspectees causes the pathways of the outcomes of interest to differ between simulations. However, they only make up a small portion of inspectees. Only 16% of inspectees are inspected in each time step (Table 4.1); within this population of inspection candidates, only a subset of inspectees are non-compliant. Depending on where these inspectees are located, the combination of peer pressure effects and their reaction to inspection and enforcement influences the behavioral pathway of the simulation. Yet, the average of 50 runs shows that RE still yields better outcomes than SE regardless of the value of $\%_{init,comp}$.

5.1.2 Varying % of Random Inspections ($\%_{rand}$)

The outcomes of interest show a larger sensitivity to the varying percentages of random inspections ($\%_{rand}$) in the mixed inspection strategy based on offense severity (MRRbOS) compared to a mixed strategy based on non-compliant record (MRRbNC); yet, this sensitivity is low in both cases (Figure 5.13). In both types of mixed random and risk-based inspection strategies (MRRbNC and MRRbOS), there is a significantly higher resulting share of compliant inspectees compared to non-compliant inspectees (Figure 5.14).

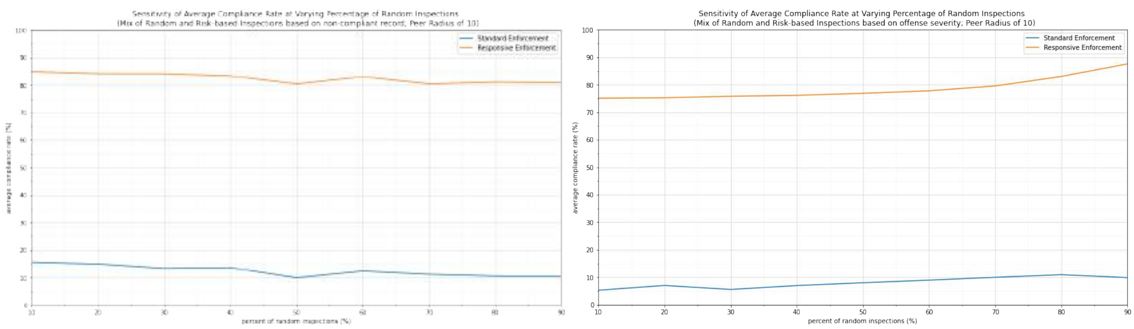

(a) MRRbNC (b) MRRbOS

Figure 5.13: Sensitivity of the average compliance rate with varying $\%_{rand}$.

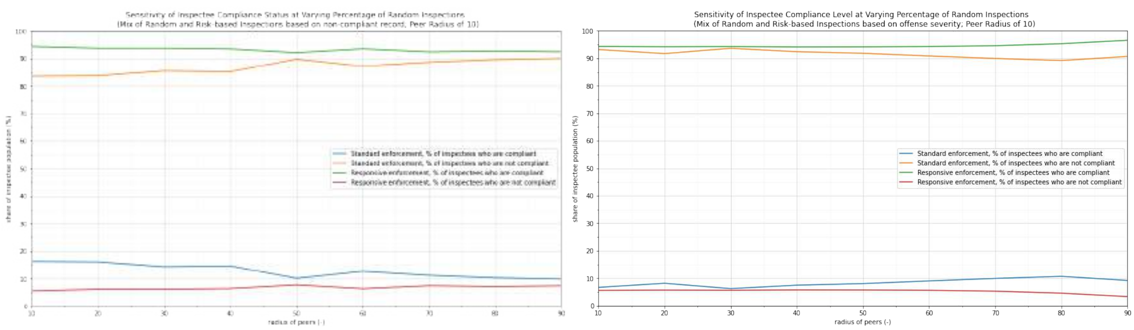

(a) MRRbNC (b) MRRbOS

Figure 5.14: Sensitivity of the compliance status breakdown with varying $\%_{rand}$.

5.1.2.1 Mixed Random & Risk-based Inspections based on Non-compliant Record (MRRbNC)

Across all values of $\%_{rand}$, the average compliance rate is higher for RE than SE (Figure 5.13a). In addition, the proportion of the inspectees who are compliant is significantly higher with RE than with SE (Figure 5.14a). With SE, the inspectee population is majority conscious violators regardless of the value of $\%_{rand}$ (Figure 5.15). As $\%_{rand}$ increases, the share of conscious violators increases. This is expected because a higher share of random inspections leaves little room for risk-based inspections. Risk-based inspections target the inspectees with the highest non-compliant record; therefore, there is a higher chance of catching non-compliant inspectees and conducting enforcement interventions that nudge the inspectee population towards higher compliance rates with a lower $\%_{rand}$.

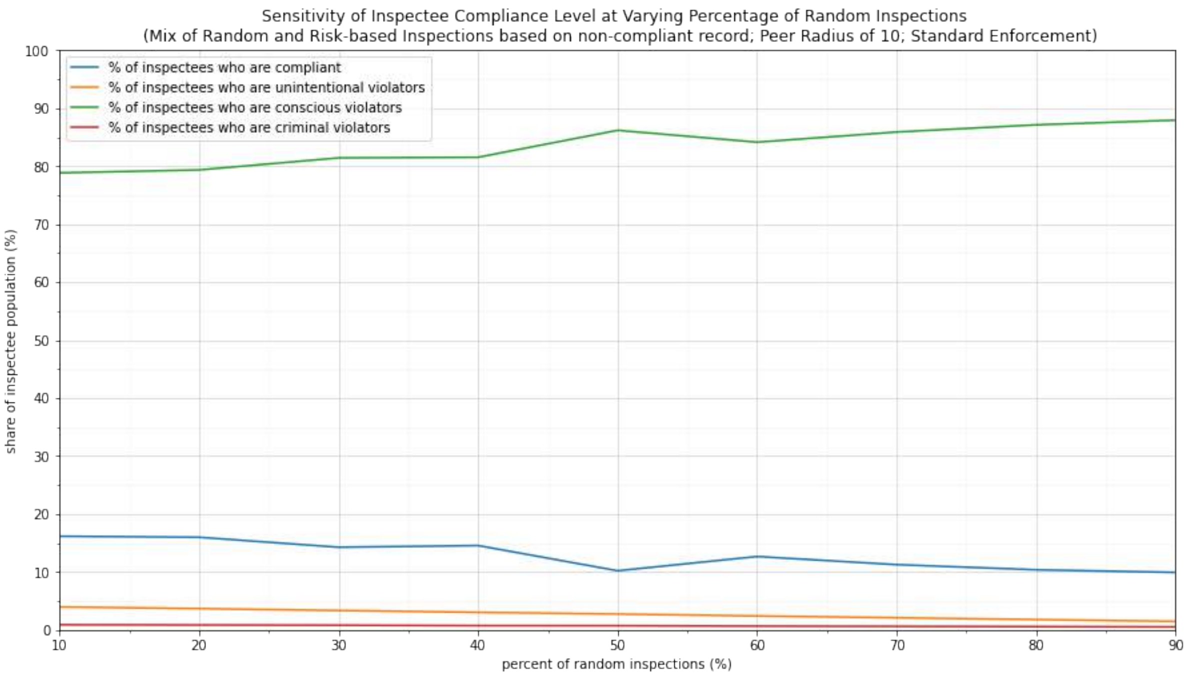

Figure 5.15: Sensitivity of the compliance level breakdown with varying $\%_{rand}$ (SE, MRRbNC).

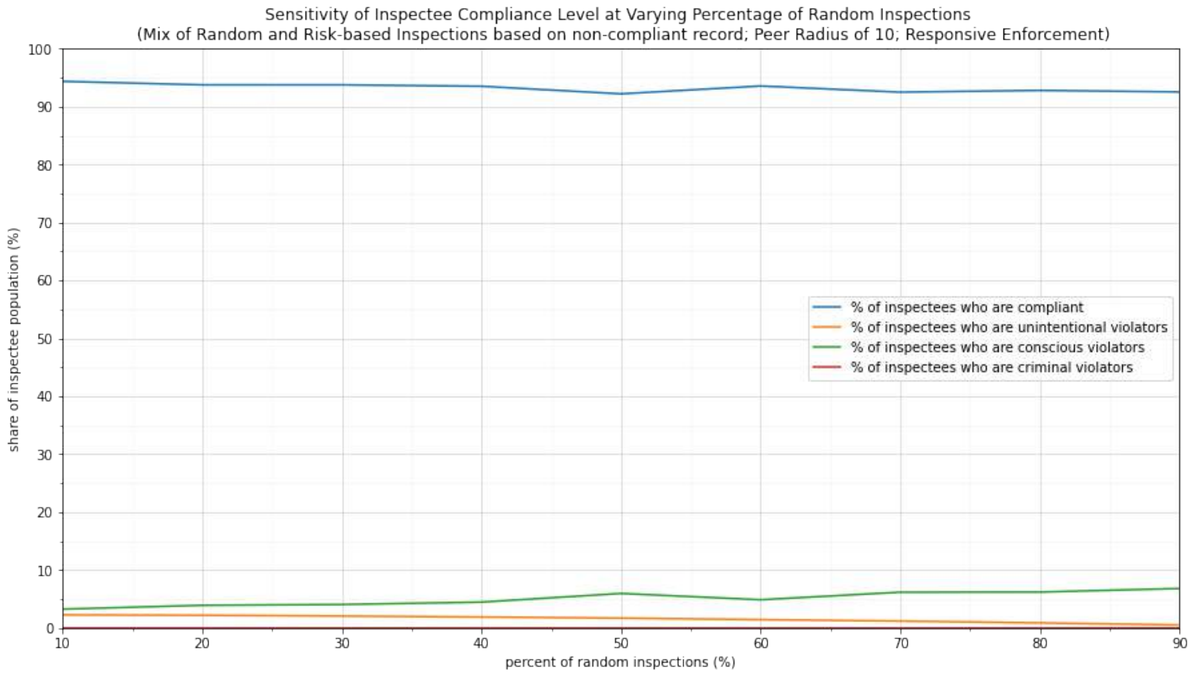

Figure 5.16: Sensitivity of the compliance level breakdown with varying $\%_{rand}$ (RE; MRRbNC).

Similarly, lower percentages of random inspections yield better outcomes than higher percentages with RE (Figure 5.16). However, with RE, the varying levels of enforcement severity render higher level offenders unable to absorb the enforcement intervention. This means more non-compliant inspectees become compliant over time, leading to a consistent upward trend of the average compliance rate and share of compliant inspectees for RE. With SE, the pathways of the outcomes of interest are more sensitive to $\%_{rand}$ (Figure 5.17). The highest compliance rates occur at 10% for both SE and RE (see Figures A.13, A.14, A.15, and A.16 for more detail).

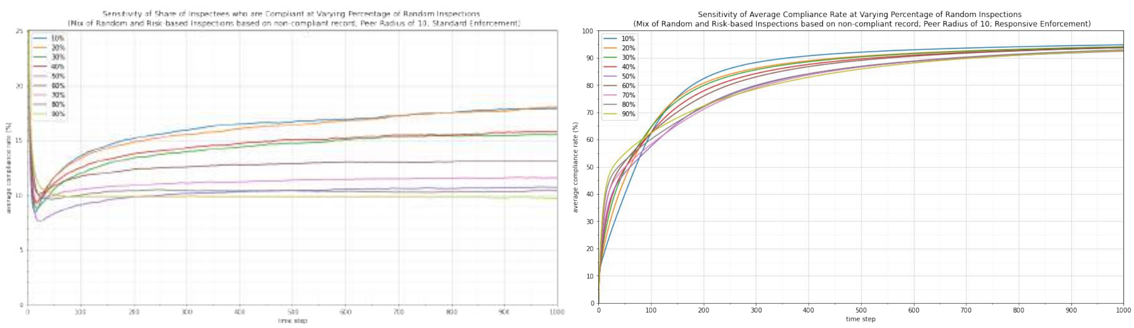

(a) Standard enforcement. (b) Responsive enforcement.

Figure 5.17: Share of compliant inspectees with varying $\%_{rand}$ (MRRbNC).

Path Dependency

The variation in the average compliance rate is much higher with RE across all values of $\%_{rand}$ than with SE (Figure 5.18). However, the variation of the share of compliant inspectees is lower with RE than with SE (Figure 5.19). This shows that more inspectees are changing their compliance level with RE compared to SE.

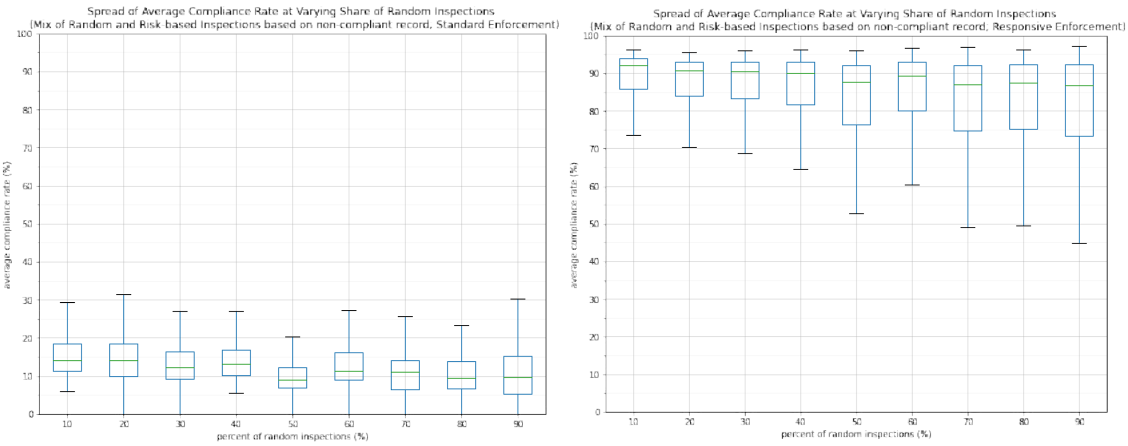

(a) Standard enforcement.

(b) Responsive enforcement.

Figure 5.18: Spread of the average compliance rate with varying $\%_{rand}$ (MRRbNC).

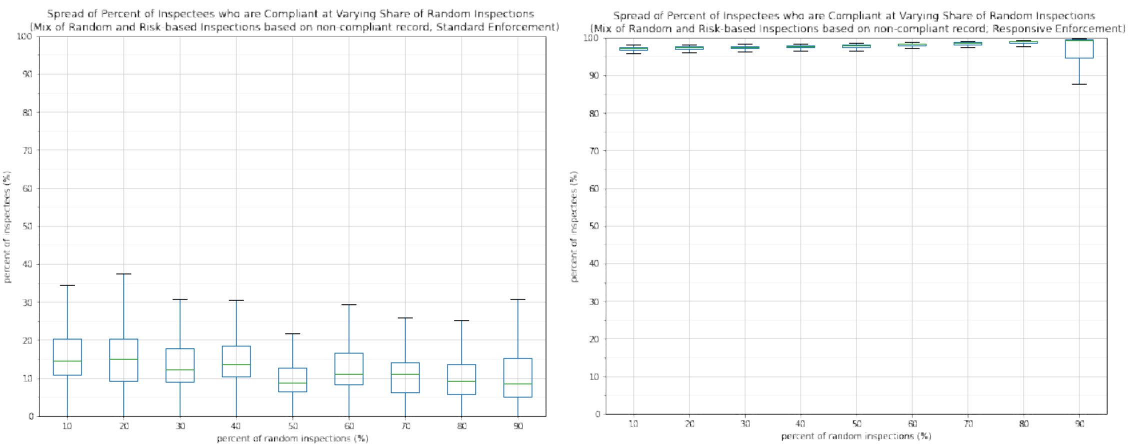

(a) Standard enforcement.

(b) Responsive enforcement.

Figure 5.19: Spread of the compliant inspectee share with varying $\%_{rand}$ (MRRbNC).

As $\%_{rand}$ increases, the path dependency becomes stronger and variation between runs widens with SE (Figures A.17 and A.18). A random selection of inspection candidates means that inspectors could be inspecting compliant agents; in this case, the opportunity to catch a non-compliant inspectee and influence their behavior is lost. As the percentage of randomly selected inspection candidates increases, this opportunity cost also increases. Yet, the difference in the outcomes between the varied values of $\%_{rand}$ is similar and within the standard deviations of each $\%_{rand}$ value (Figures 5.18a and 5.19a).

A similar pattern occurs with RE, but with less variation and path dependency (Figures A.19 and A.20). At lower values of $\%_{rand}$, there is less variation in the results because more inspection candidates are chosen based on non-compliant record than at random. The targeting of inspection candidates based on their historical non-compliance ensures that repeat offenders are selected, increasing the chances that the inspectee will change its behavior after the intervention.

In sum, the lower the $\%_{rand}$, the better the outcomes of interest. At lower values of $\%_{rand}$, there is a lower opportunity cost of not inspecting non-compliant actors. Therefore, the outcomes of interest fare better than at higher $\%_{rand}$ values where there could potentially be a larger portion of compliant inspectees that are selected for inspection. In addition, the higher values of $\%_{rand}$ have greater path dependency and variation in outcomes due to the large number of randomly selected inspectees.

5.1.2.2 Mixed Random & Risk-based Inspections based on Offense Severity (MRRbOS)

RE yields better outcomes compared to SE at all values of $\%_{rand}$ when a MRRbOS strategy is used (Figures 5.20 and 5.21). Contrary to the MRRbNC strategy, targeting inspections based on offense severity results in a higher average compliance rate (Figures 5.20). However, there is little sensitivity of the share of compliant inspectees left in the population with varying values of $\%_{rand}$ (Figures 5.21 and 5.22). While the risk-based inspections target the highest level offenders, SE is not sufficient to deter conscious and criminal offenders from continuing to violate; unintentional violators are less likely to be able to absorb the severity of SE so they become more compliant. As explained previously, RE forces higher level offenders to comply as they cannot absorb the higher level of enforcement severity.

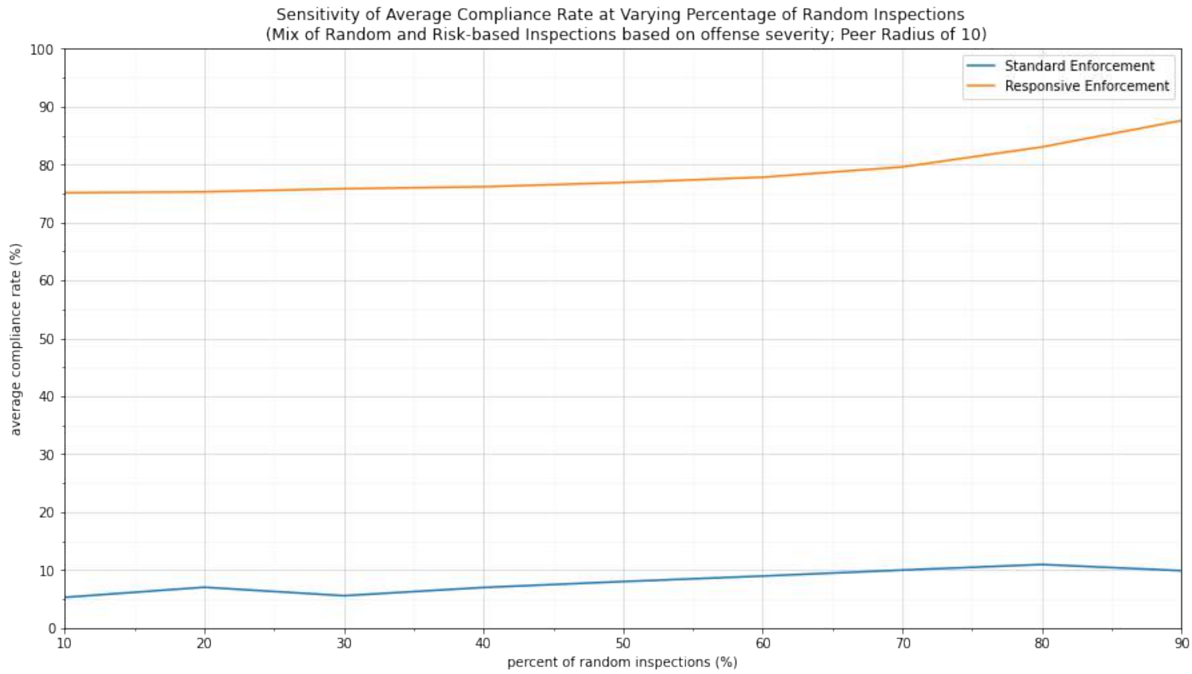

Figure 5.20: Sensitivity of the average compliance rate with varying $\%_{rand}$ (MRR-bOS).

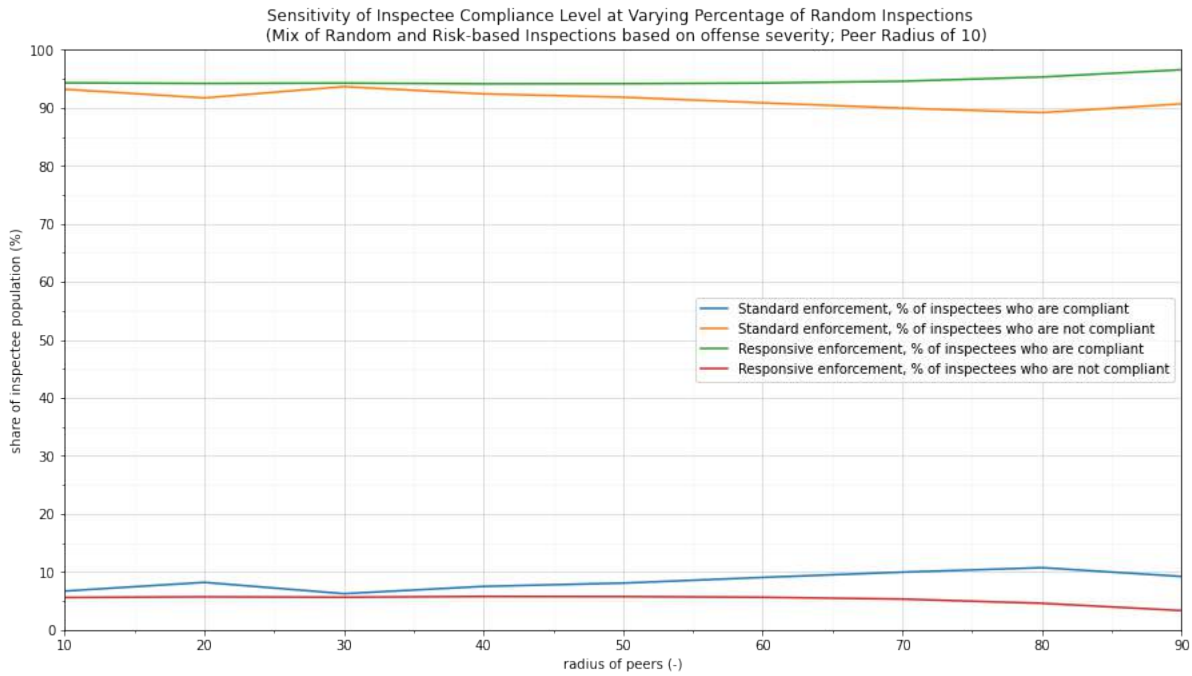

Figure 5.21: Sensitivity of the compliance status breakdown with varying $\%_{rand}$ (MRRbOS).

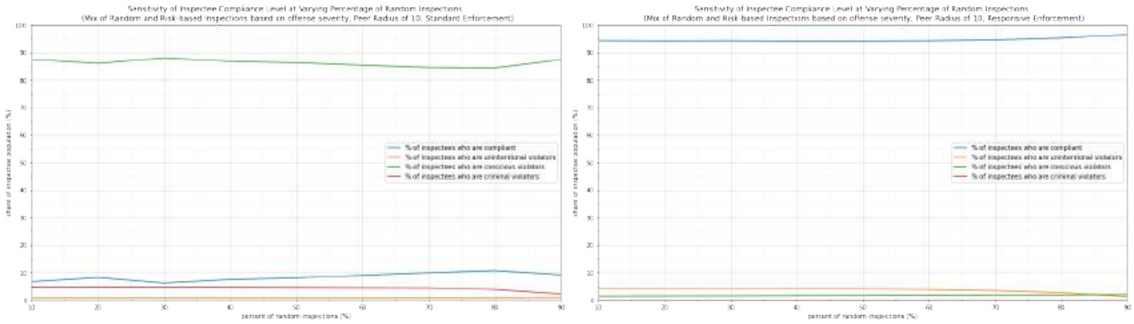

(a) Standard enforcement. (b) Responsive enforcement.

Figure 5.22: Sensitivity of the compliance level breakdown with varying $\%_{rand}$ (MR-RbOS).

The behavior of the outcomes of interest over time is similar at varying percentages of random inspections, indicating little sensitivity to $\%_{rand}$ (Figure 5.23; also see Figures A.21 and A.22). However, with SE, the difference between the outcomes of interest at each varied value is wider than with RE.

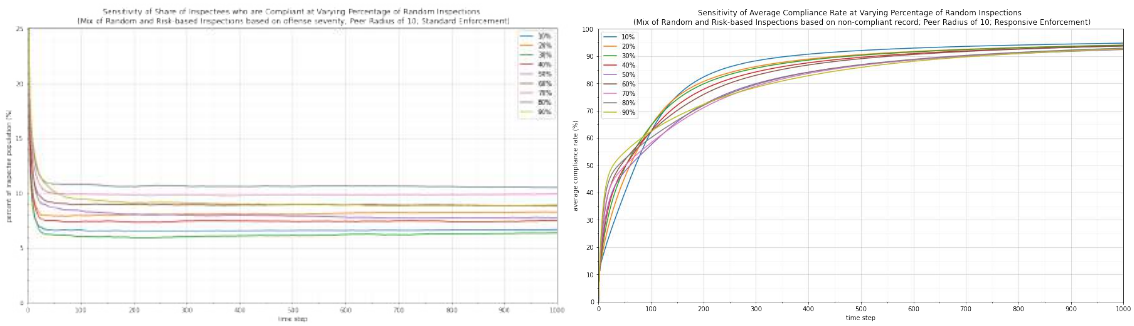

(a) Standard enforcement. (b) Responsive enforcement.

Figure 5.23: Share of compliant inspectees over time with varying $\%_{rand}$ (MRRbOS).

Path Dependency

There is a wider variation and stronger path dependency with SE compared to RE (Figures 5.24 and 5.25; also see Figures A.26, A.27, A.28, and A.29). With SE, the path dependency strengthens as the $\%_{rand}$ increases (Figures A.26 and A.27). Compared to the MRRbNC strategy, the MRRbOS strategy targets higher level offenders more precisely; therefore, it requires relatively less risk-based inspections to improve the average compliance rate holistically. The simulations with SE are highly path dependent, indicating that the location of inspectees highly influences the macro-level behavior that is observed. Additionally, there is more path dependency as $\%_{rand}$ increases for RE (Figure A.28). The higher share of random inspections allows for more variety of actors to be inspected. Yet, the path dependency is limited due to the effectiveness of RE to influence inspectees to become compliant.

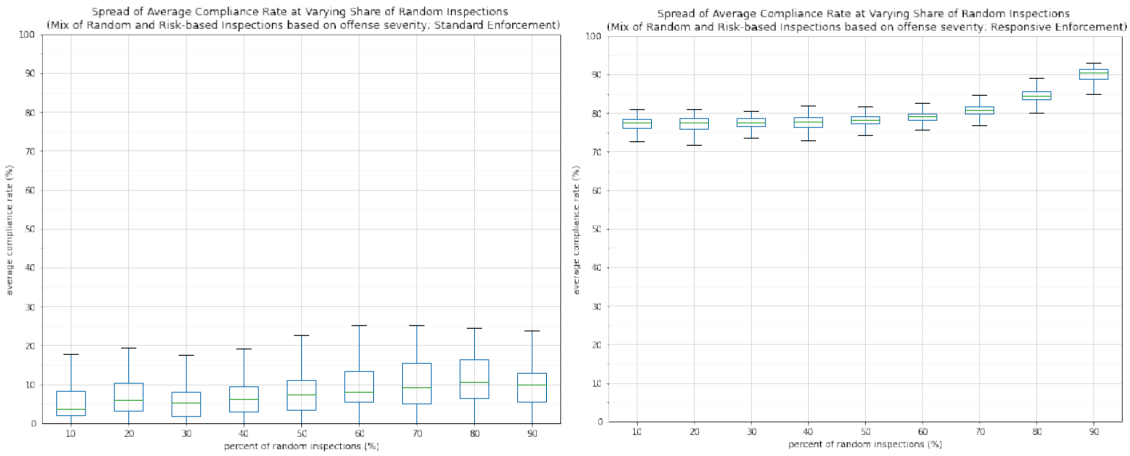

(a) Standard enforcement. (b) Responsive enforcement.

Figure 5.24: Spread of average compliance rate with varying $\%_{rand}$ (MRRbOS).

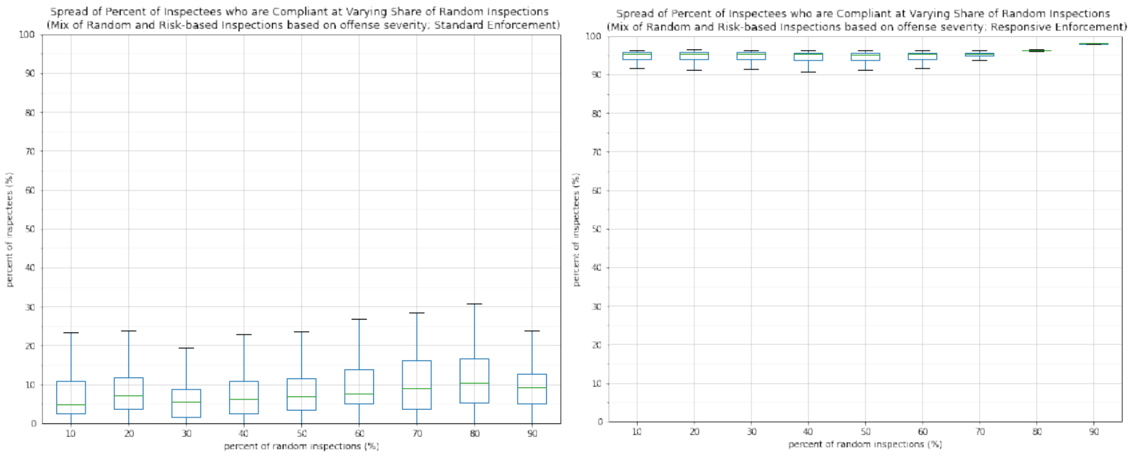

(a) Standard enforcement. (b) Responsive enforcement.

Figure 5.25: Spread of the compliant inspectee share with varying $\%_{rand}$ (MRRbOS).

To summarize, the outcomes of interest have little sensitivity to $\%_{rand}$ when a MRRbOS inspection strategy is used. Targeting higher level offenders helps to bring down the share of non-compliant inspectees with little path dependency, particularly with RE. SE is not sufficient to deter higher level offenders, though effective in influencing unintentional violators to become compliant. It is worth noting that the $\%_{rand}$ parameter has been used in previous ILT inspection models, such as Knol's game theoretic ABM (see Section 2.1.1), to introduce learning characteristics to inspector agents. These previous models assume that an inspector agent has no insight into the characteristics of an inspectee until they are inspected; therefore, having a certain percentage of random inspections allows for the inspectors to collect data on previously unknown inspectees. However, the PABMI uses empirical data to assign the historical non-compliance record to each inspectee at the initialization of the model. The risk-based inspection candidates are selected based on a statistical distribution of historical compliance record that is defined at setup. Therefore, the $\%_{rand}$ parameter in the PABMI simply widens the pool of inspection candidates beyond those targeted by risk-based inspections.

5.1.3 Varying Logit Parameters (μ and σ)

The logit μ and σ parameters indicate the mean and standard deviation of the distribution of absorbance capacity, respectively (for details on the population distribution by absorbance capacity, see Figures A.30 through A.39). A high (low) logit μ parameter indicates that the inspectee population has a high (low) tolerance for enforcement on average. A high (low) logit σ parameter indicates the inspectees differ (are similar) in their capacity to absorb enforcement.

The results of the sensitivity analysis show that with SE, lower values of μ yield higher compliance rates (Figure 5.26a; also see Figures A.40a and A.41a). Lower values of μ mean that a large number of inspectees do not have the capacity to absorb the severity of SE, which is a fixed value regardless of the offense severity (see Figure 3.2). Larger values of σ mean that there is a wider range of absorbance capacities among the inspectees, making the overall population more resilient to SE. With RE, a logit μ of less than or equal to 0.7 means that a majority of the inspectees are unable to absorb the severity of RE even with a large σ (Figures 5.26b; also see Figures A.40b and A.41b). This is structurally invalid, as the population of inspectees becomes exclusively compliant too quickly.

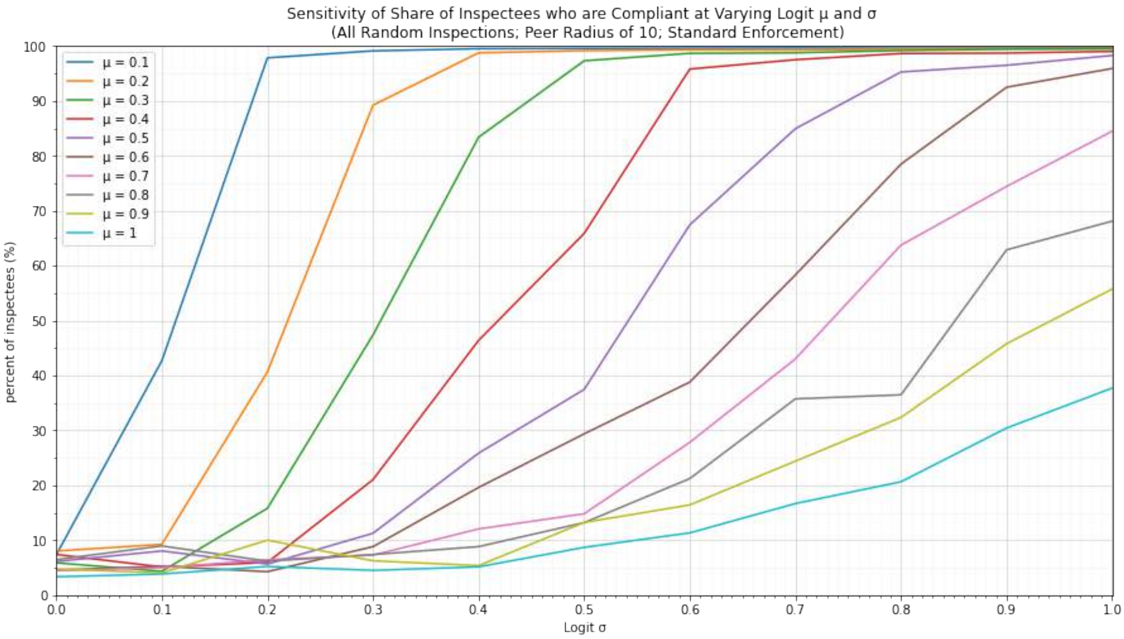

(a) Standard enforcement

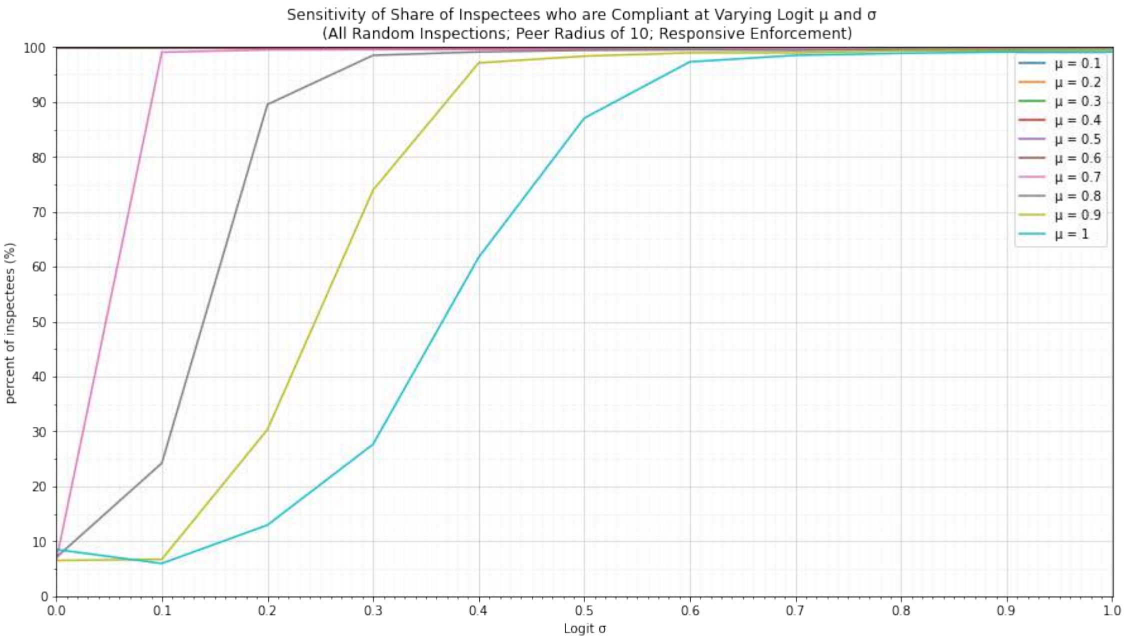

(b) Responsive enforcement.

Figure 5.26: Sensitivity of the share of compliant inspectees with varying μ & σ.

The structurally valid values of the logit parameters are shown in Table 5.2. The sensitivity analysis results show that if inspectees react to enforcement (see Section 3.4), there must be a reasonable assumption about their capacity to withstand the enforcement intervention. Without this, the PABMI yields results that are implausible. However, the inspectees' reaction to enforcement can also be switched off in the PABMI; in this case, the logit parameters are obsolete.

	Valid range of Logit σ Values	
μ	Standard Enforcement	Responsive Enforcement
0.1	0.11 - 0.2	None
0.2	0.22 - 0.3	None
0.3	0.31 - 0.4	None
0.4	0.42 - 0.6	None
0.5	0.54 - 0.8	None
0.6	0.66 - 1.0	None
0.7	0.73 - 1.0	None
0.8	0.85 - 1.0	0.14 - 0.3
0.9	0.94 - 1.0	0.24 - 0.4
1	None	0.37 - 0.6

Table 5.2: Structurally valid values of μ and σ.

5.1.3.1 Varying Logit Sigma (σ)

As previously mentioned, a wider distribution (σ) of the inspectees across all possible absorbance capacities yields better overall outcomes for both SE and RE (Figures 5.27, A.44, A.45). However, RE is more effective for increasing compliance, as more severe enforcement interventions that come with RE make a greater portion of non-compliant inspectees unable to absorb the enforcement severity.

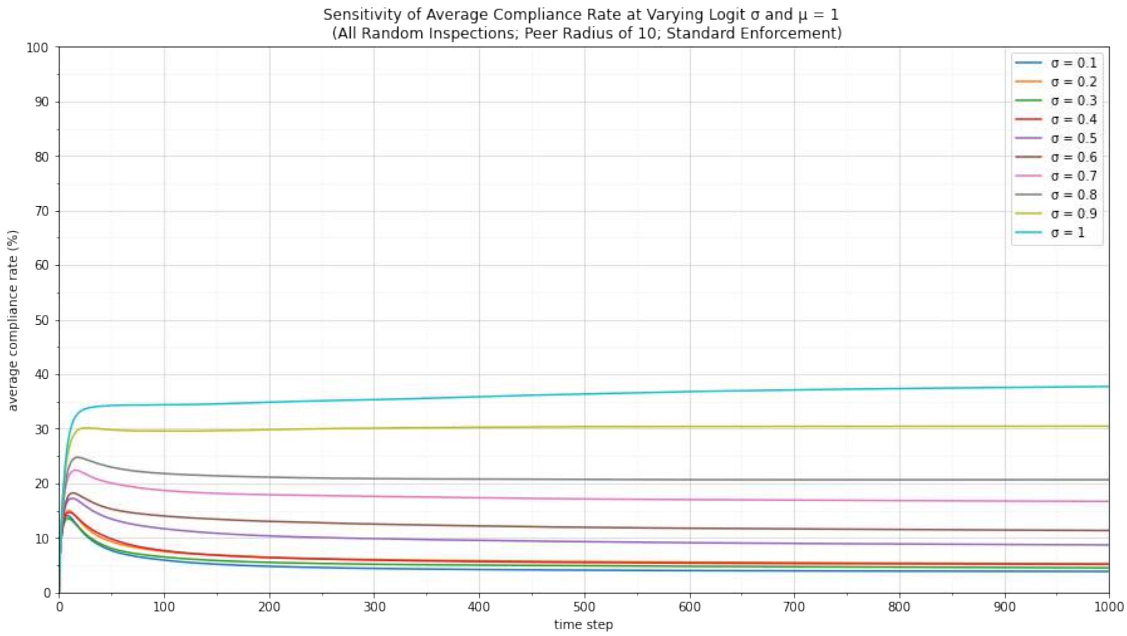

(a) Standard enforcement

Figure 5.27: Average compliance rate results with varying σ & μ=1.

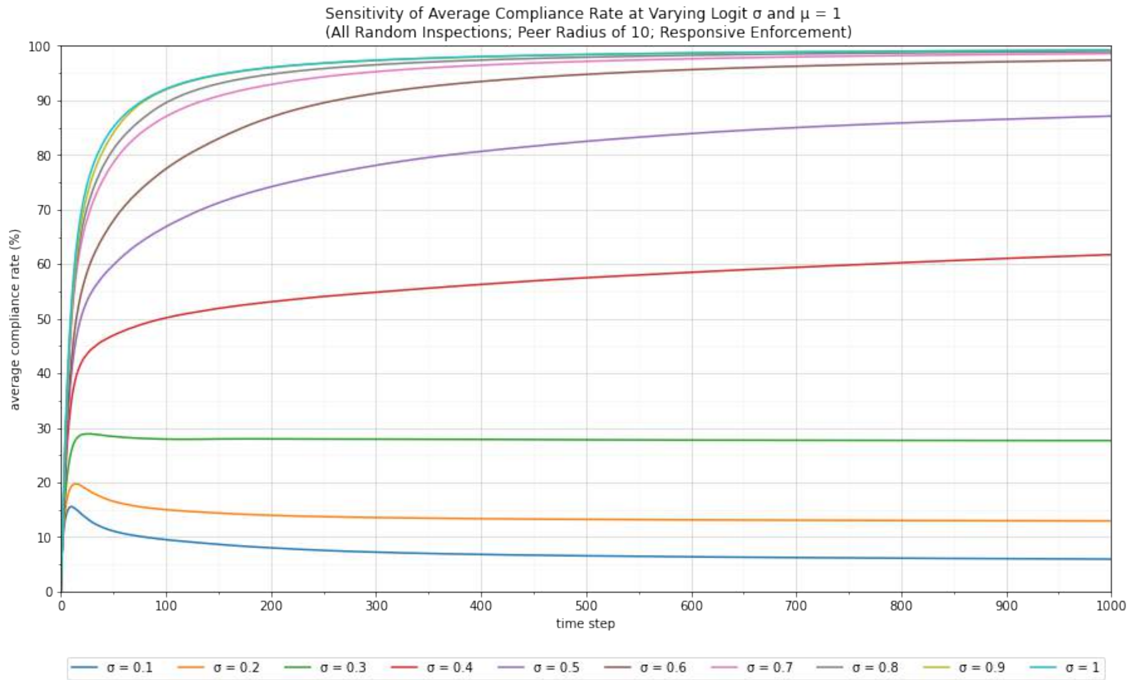

(b) Responsive enforcement.

Figure 5.27: Average compliance rate results with varying σ & μ=1 (continued).

Path Dependency

As σ increases, the variation in the outcomes of interest increases as expected (Figures 5.28 and 5.29). There is a wider variation with RE, showing that there are more inspectees that change their behavior than with SE. A large standard deviation, particularly 0.6 and above, shows high variation where the outcomes of interest are strongly path dependent (Figure 5.29b).

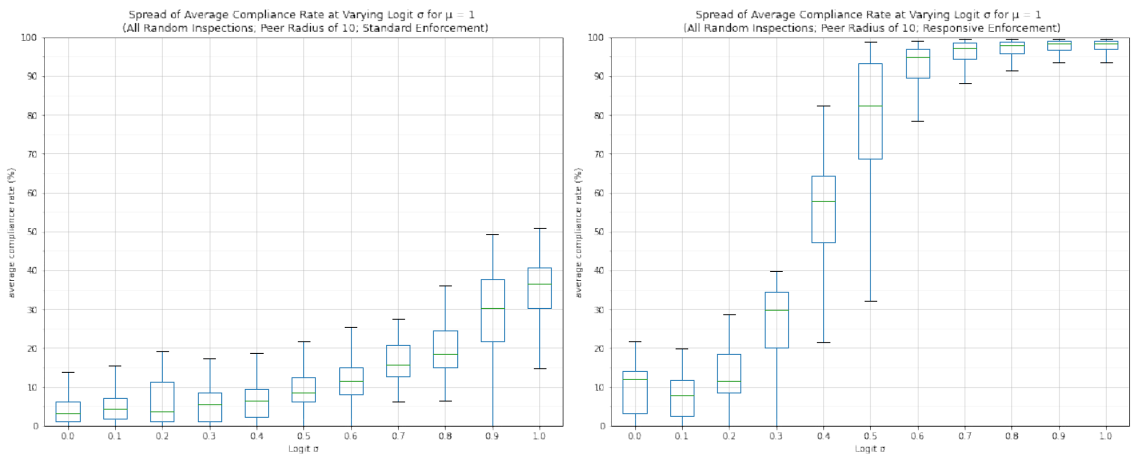

(a) Standard enforcement. (b) Responsive enforcement.

Figure 5.28: Spread of the average compliance rate with varying σ & μ=1.

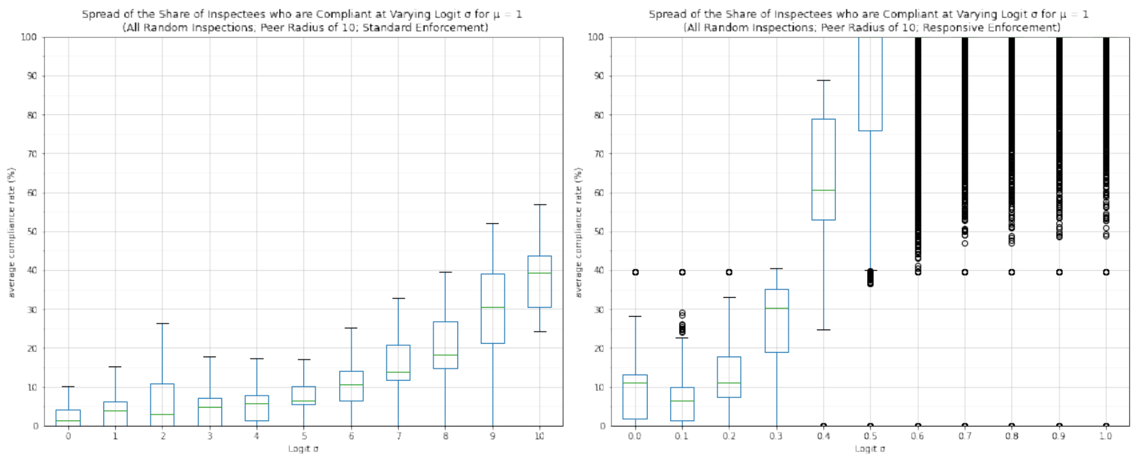

(a) Standard enforcement.

(b) Responsive enforcement.

Figure 5.29: Spread of the compliant inspectee share with varying σ & μ=1.

With SE, there is stronger path dependency as the standard deviation (σ) increases (Figures A.48 and A.49). This shows that a wider range of absorbance capacities across the inspectee population causes more changes in behavior, leading to more varied pathways depending on the location of inspectees. With RE, the highest path dependency occurs between logit μ values of 0.4 and 0.5 (Figures A.50 and A.51). A low σ indicates a low variation in absorbance capacities among inspectees. Therefore, all inspectees with an offense severity that is less than the absorbance capacity can continue violating, causing a plateau in compliance outcomes. Large values of σ mean there is a wider range of absorbance capacities; those with lower absorbance capacities will quickly become compliant. Coupled with the higher upper boundary of the severity of RE and peer pressure, the population quickly becomes majority compliant.

5.1.3.2 Varying Logit Mu (μ)

As the logit μ increases, the compliance outcomes worsen (Figures 5.30, A.52, and A.53). Larger μ values indicate higher absorbance capacities, so inspectees are able to absorb the enforcement interventions at higher severity levels and continue non-compliant behavior.

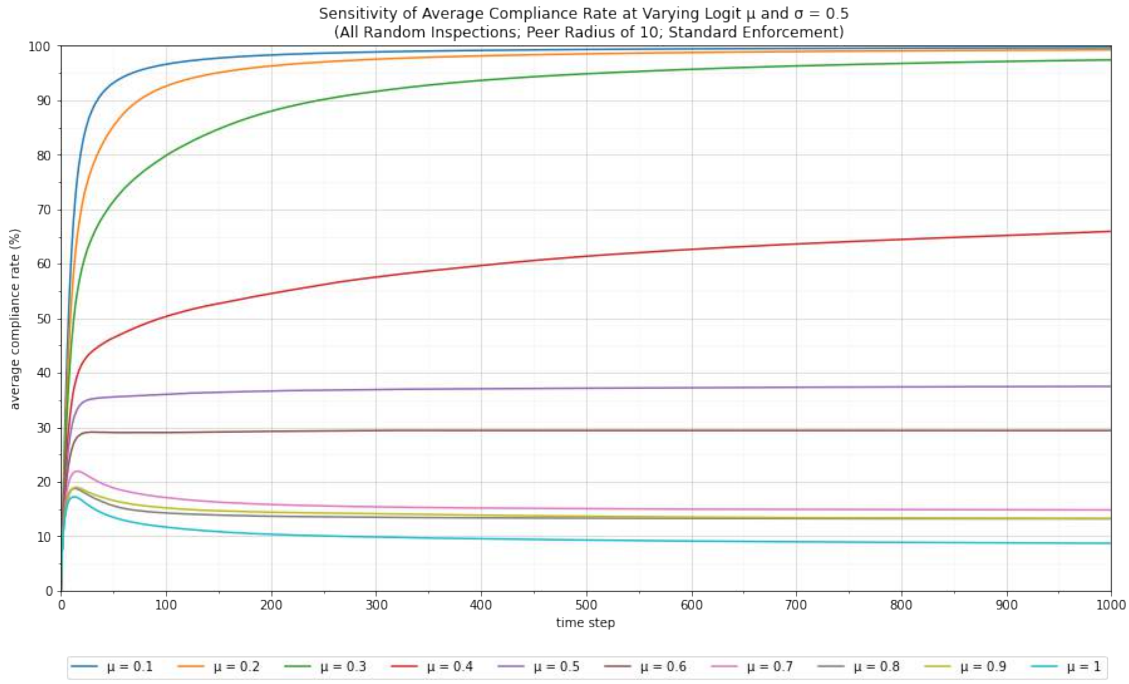

(a) Standard enforcement

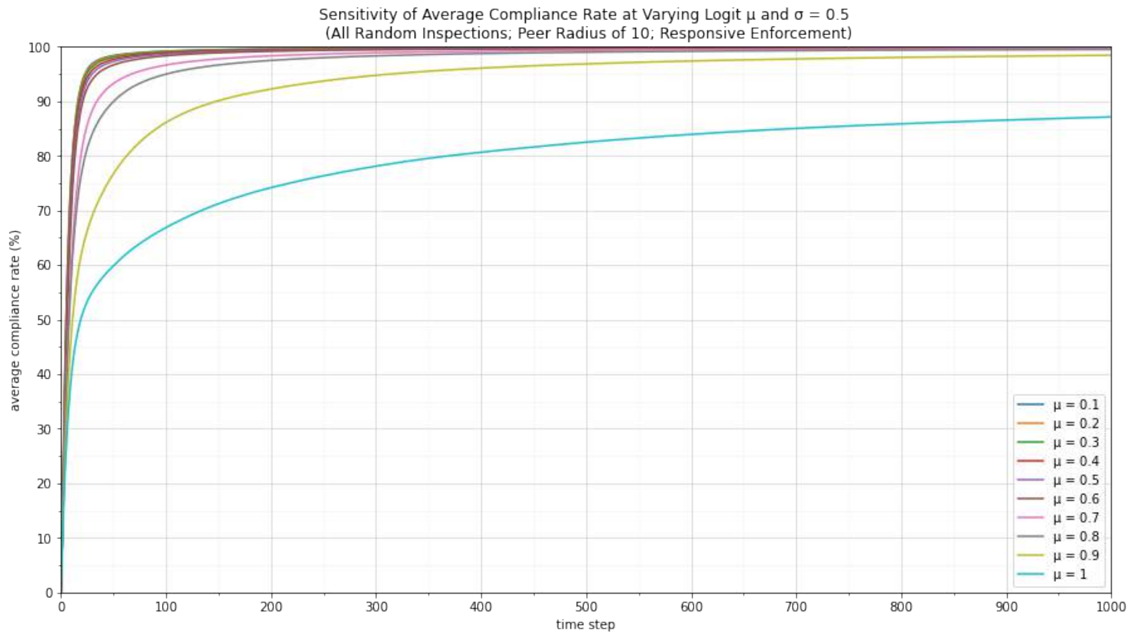

(b) Responsive enforcement.

Figure 5.30: Average compliance rate results with varying σ & μ=1.

Path Dependency

With both SE and RE, there is a large amount of variation in the outcomes of interest, indicating high path dependency (Figures 5.31 and 5.32). With SE, μ values less than or equal to 0.3 mean that few inspected agents can withstand the impact of enforcement (Figures A.56 and A.57); therefore, those that are inspected become compliant quickly, increasing the percentage of compliant inspectees. As μ increases, more inspectees are able to withstand the impact of enforcement, making the model outcomes less path dependent.

With RE, only high values of μ (i.e., μ of 0.9 and 1) show path dependency (Figure A.58 and A.59). This shows that with less variation in the absorbance capacities, many inspectees cannot withstand the impact of enforcement and become compliant quickly. RE demands high absorbance capacities from inspectees if they want to continue violating. The path dependency in the graphs are caused by the variation between the absorbance capacities of randomly selected inspection candidates.

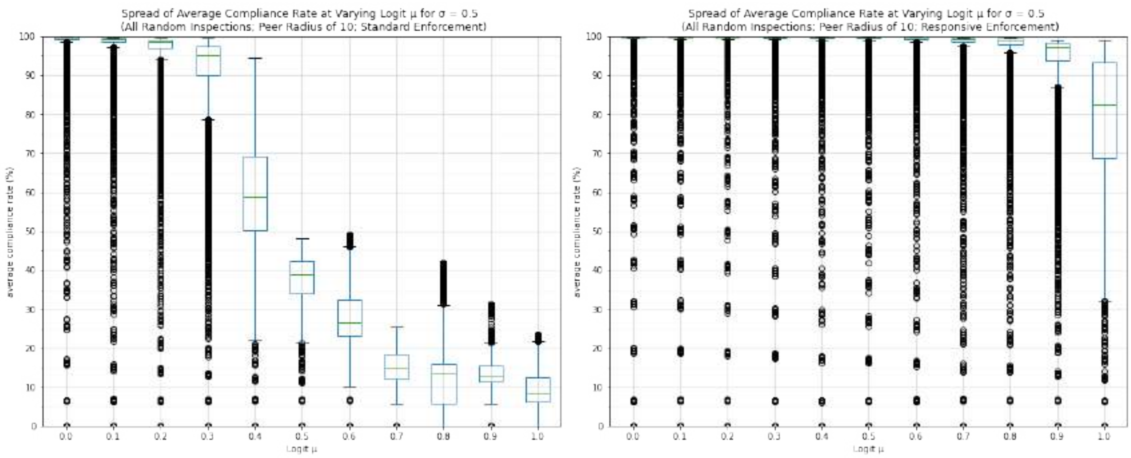

(a) Standard enforcement.

(b) Responsive enforcement.

Figure 5.31: Spread of the average compliance rate at varying μ & σ=0.5.

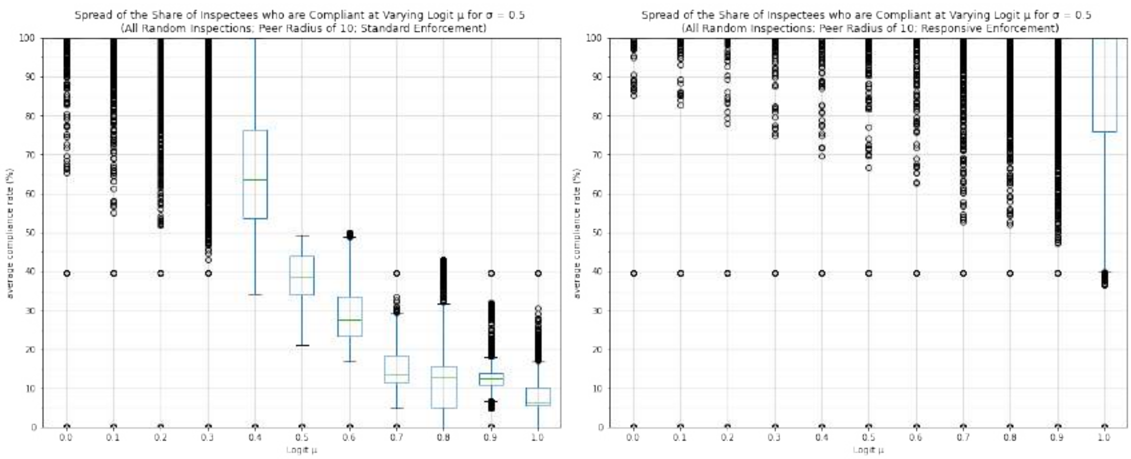

(a) Standard enforcement.

(b) Responsive enforcement.

Figure 5.32: Spread of the compliant inspectee share with varying μ & σ=0.5.

5.1.4 Varying Radius of Peers (r_{peers})

In an AR inspection strategy, the outcomes of interest are highly sensitive to the radius of peers (Figure 5.33 and 5.34). The highest average compliance rate is achieved when the radius of peers (r_{peers}) is 10 and RE is conducted (Figure 5.33 and 5.34). RE consistently yields a higher average compliance rate regardless of value of r_{peers} compared to SE (Figure 5.33).

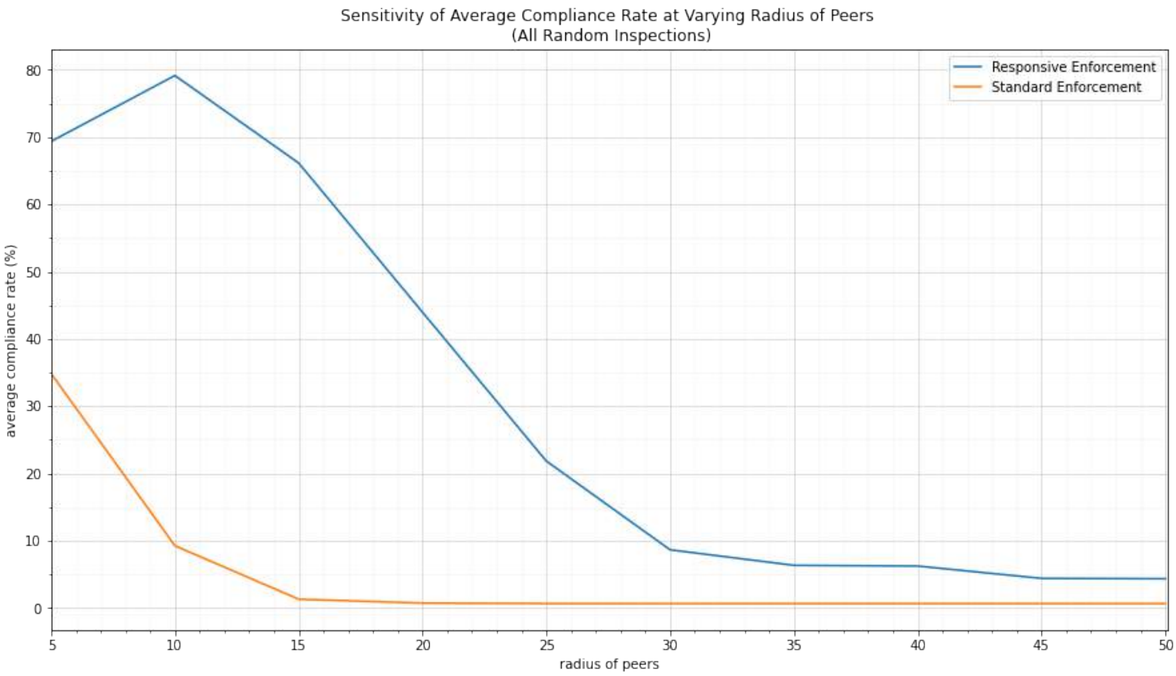

Figure 5.33: Sensitivity of the average compliance rate with varying r_{peers}.

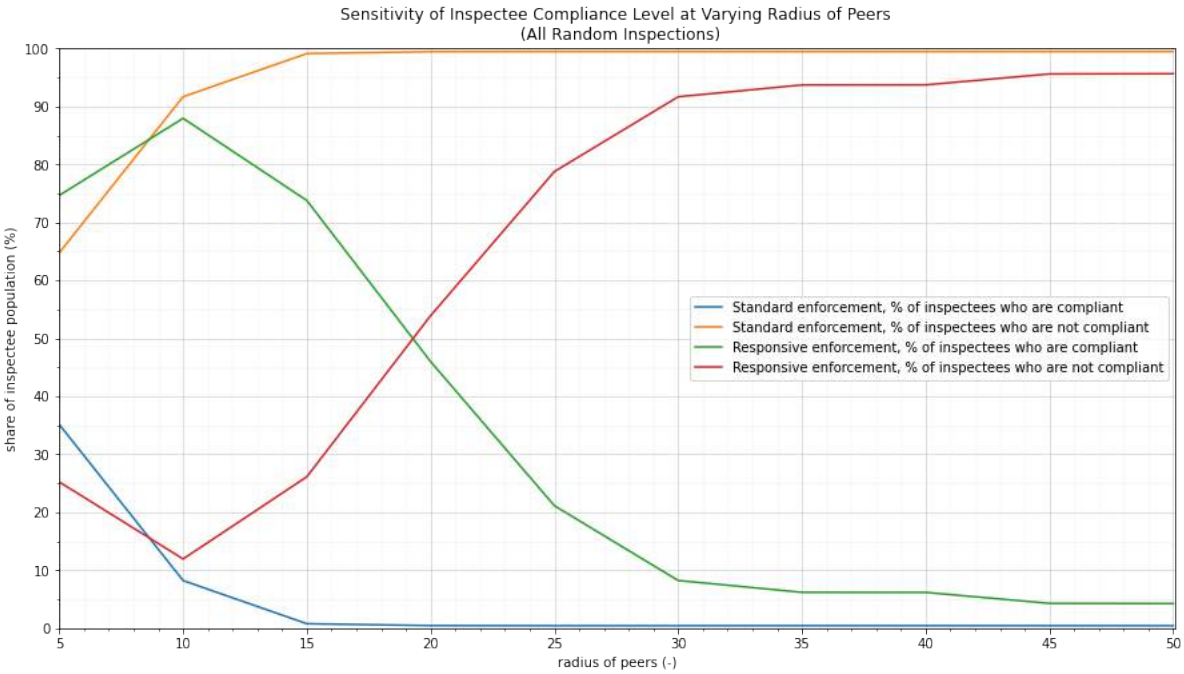

Figure 5.34: Sensitivity of the compliance status breakdown with varying r_{peers}.

Standard Enforcement

Across all values of r_{peers}, the inspectee population consists of an overwhelming majority of conscious violators (Figure 5.35). Given the default values of the model parameters (Table 4.1), conscious violators make up the largest majority of the inspectee population at the start of the simulation. This means that as r_{peers} increases, the likelihood that an inspectee will have more non-compliant than compliant neighbors increases (see Section 3.4.3). In addition, the distribution of the absorbance capacity with logit normal μ an σ default values of 1.0 and 0.5, respectively, produces a population of inspectees where anywhere between 93% to 98% of the inspectees have an absorbance capacity greater than 0.5 (see Figure A.39 and Appendix A.1.3 for details). This means that an overwhelming majority of inspectees can absorb the standard severity of enforcement, so they continue violating at their current compliance level (see Figure 3.2). The effect of peer pressure coupled with the ability of inspectees to cope with the severity of enforcement produce a population of inspectees where the behavior of conscious violators is most often adopted during SE (Figure 5.35).

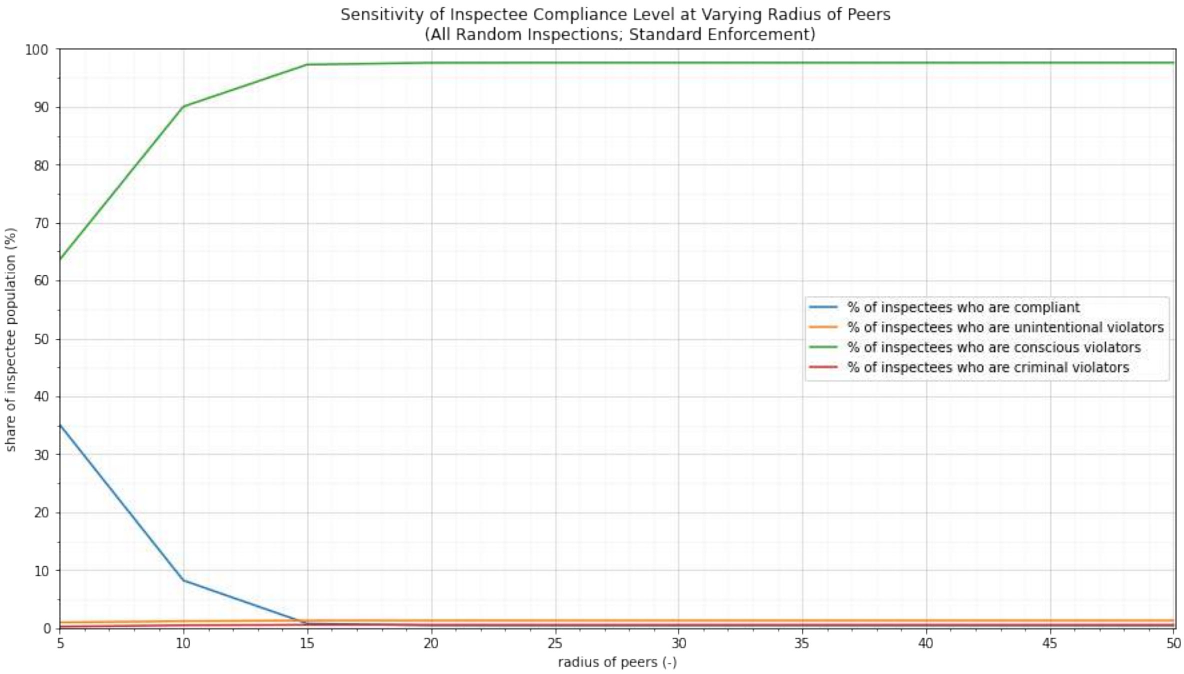

Figure 5.35: Sensitivity of the compliance level breakdown at varying r_{peers} (SE).

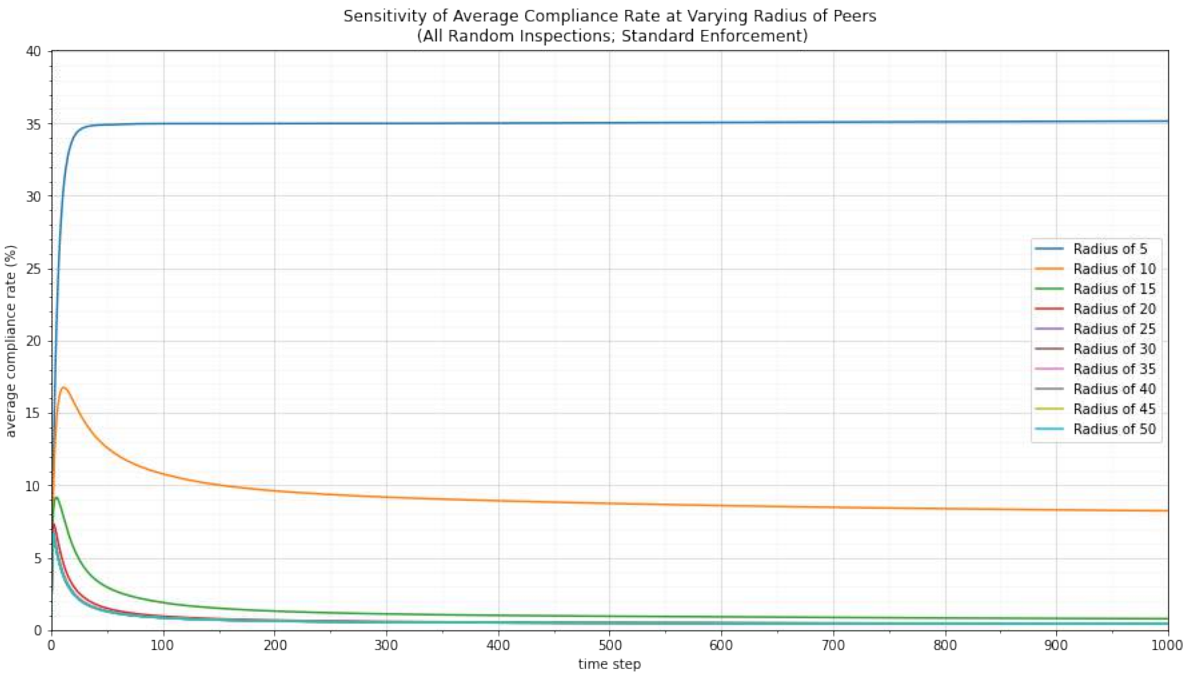

Figure 5.36: Average compliance rate results with varying r_{peers} (SE).

Smaller values of r_{peers} cause the outcomes of interest to reach steady state quicker than with larger values of peer radius (Figures 5.36, 5.37; also see Figures A.61 and A.60). This is expected, as a smaller r_{peers} means there are fewer agents an inspectee considers a peer. Therefore, it is influenced by fewer agents. The behavior of a smaller circle of peers becomes homogenized quicker than with a larger one.

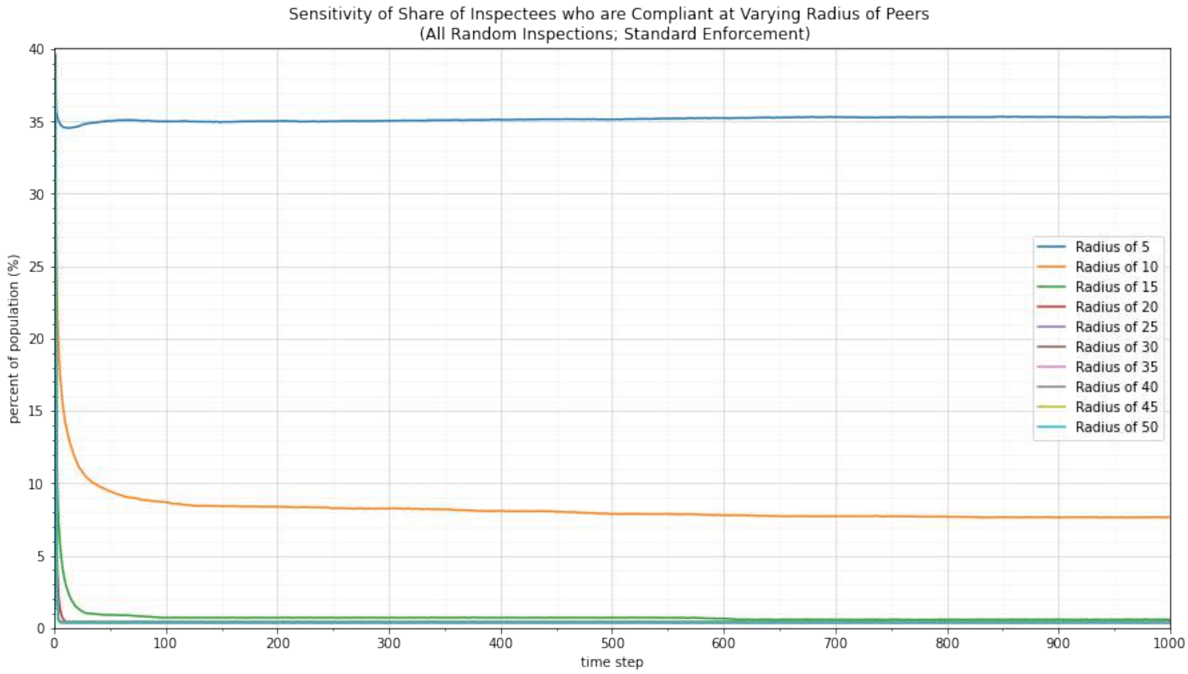

Figure 5.37: Share of compliant inspectees over time with varying r_{peers} (SE).

With an AR inspection strategy and SE, the upper limit of r_{peers} is five. If r_{peers} is set to 10, more than 90% of inspectees become non-compliant in a short amount of time; if r_{peers} is 15 or higher, 100% of the inspectees become non-compliant within the first 30 inspection cycles (Figure A.60). This is highly unrealistic. Therefore, the effect of peer pressure should be limited to inspectees within a radius of five.

Responsive Enforcement

With RE, the number of compliant inspectees exceeds the number of non-compliant inspectees when r_{peers} is less than or equal to 15; with a radius greater than 20, the number of non-compliant inspectees exceeds the number of compliant inspectees (Figure 5.38). This shows that when RE is utilized, more conscious and criminal offenders are likely to become compliant, as they cannot absorb the severity of the enforcement. As more inspectees become compliant, the effect of peer pressure causes other inspectees to become compliant. The effect of peer pressure is stronger than the severity of enforcement at radii greater than 20 when RE is performed (Figure 5.38).

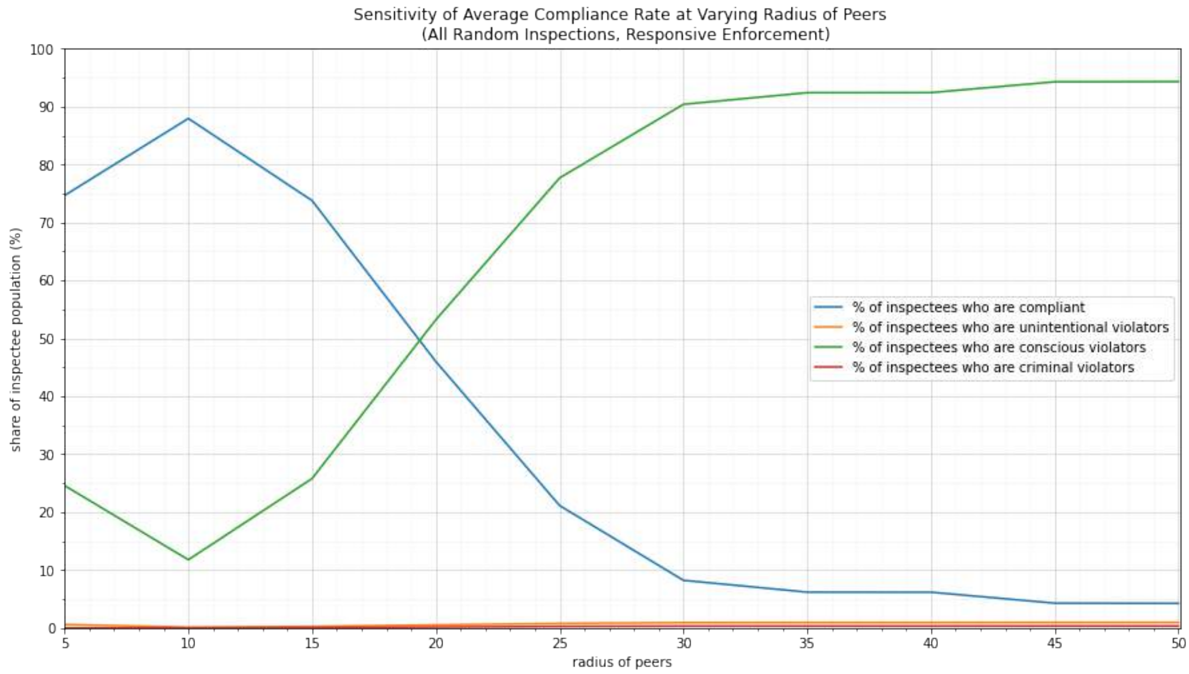

Figure 5.38: Sensitivity of the compliance level breakdown with varying r_{peers} (RE).

The optimal r_{peers} for RE is 10, as this yields the highest average compliance rate and share of compliant inspectees (Figure 5.39; also see Figure A.62). Unlike SE, a larger r_{peers} does not predictably yield lower average compliance rates or the compliant share of the inspectee population (Figures 5.39; also see Figures A.62, A.63, and A.64). This shows that RE mitigates the effect of non-compliant peer pressure, especially for conscious and criminal violators; a higher severity of enforcement for these violators renders them less capable of absorbing it.

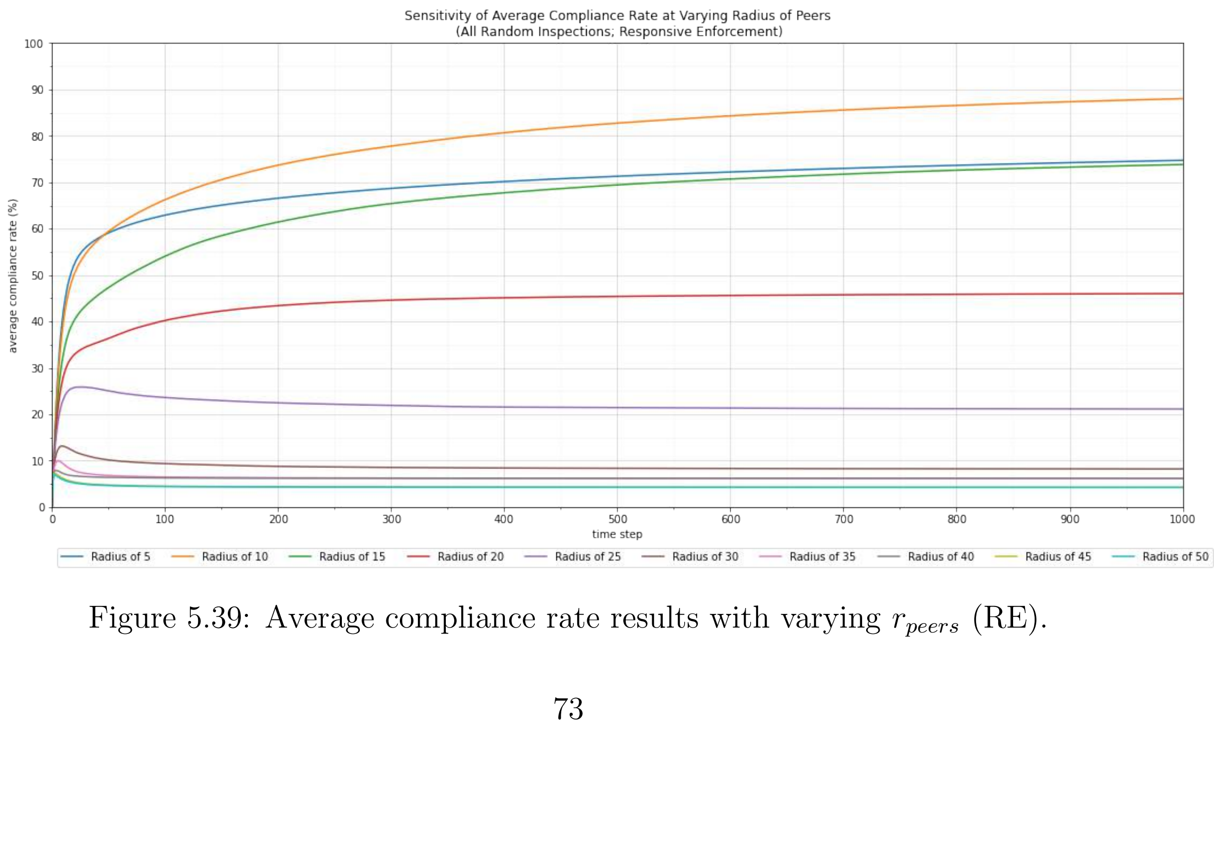

Figure 5.39: Average compliance rate results with varying r_{peers} (RE).

Only radii between five and 15 yield a non-compliant population that is less than 50% (Figure A.63). Therefore, r_{peers} should be limited to 15 or less to be valid and set to 10 to produce the optimal average compliance rate.

Path Dependency

There is wider variation in the average compliance rate with RE than SE (Figure 5.40). A couple of factors could contribute to this. First, since the inspectees are placed randomly in the model world, inspectees of various compliance statuses are congregated randomly; for each run of the model, peer pressure varies depending on where the inspectees are placed. Second, more inspectees make more drastic changes in their compliance level during RE than SE (e.g., a criminal offender becoming compliant in one time step occurs more frequently during RE than SE); this occurs because the likelihood that a conscious or criminal offender can absorb the severity of RE is lower than that of SE. As violating inspectees change their compliance level more drastically, the average compliance rate of the overall population also changes more abruptly. The share of compliant inspectees also follows this behavior (see Figures 5.41).

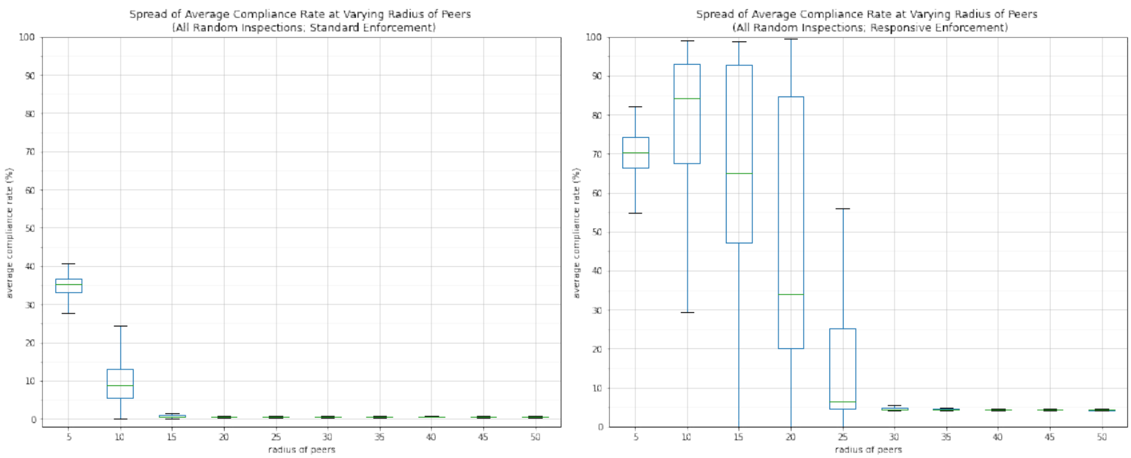

(a) Standard enforcement. (b) Responsive enforcement.

Figure 5.40: Spread of the average compliance rate with varying r_{peers}.

With SE and the r_{peers} of five, the pathway of each simulation shows similar trajectories (see Figures A.66, A.67, and A.68 for details). However, the average compliance rate and share of compliant inspectees at steady state can vary by up to 12% and 15%, respectively (see Figures A.66 and A.67 for details). This shows that the starting position of the inspectees influences the outcomes of interest within a range of variation.

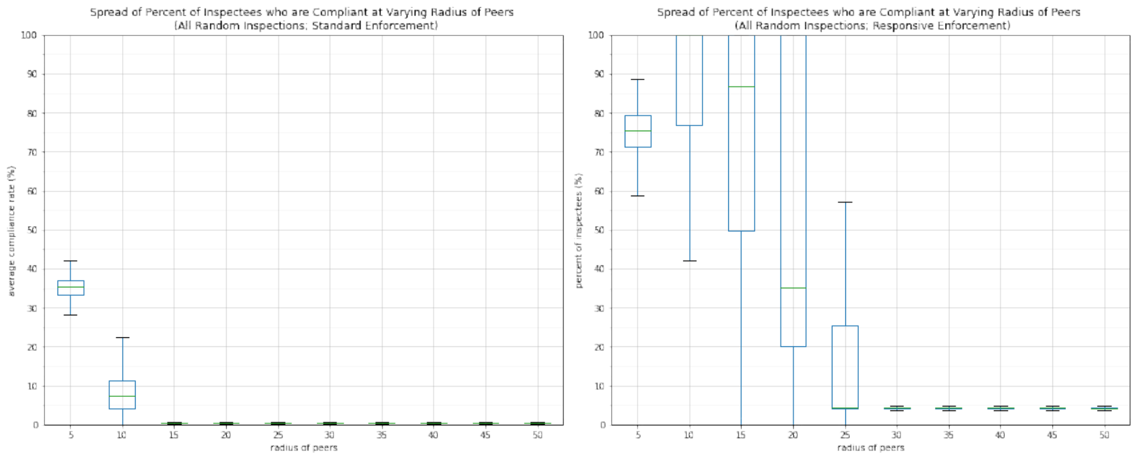

(a) Standard enforcement. (b) Responsive enforcement.

Figure 5.41: Spread of the compliant inspectee share with varying r_{peers}.

With RE, the trajectories of the outcomes of interest are similar for each run if the r_{peers} is five (Figures 5.42). However, if the r_{peers} equals 10 or 15, the path dependency increases drastically (Figure 5.42). This shows that as the radius of peers widens, the variation in the effect of peer pressure becomes stronger; in this case, the location of the inspectees in relation to its peers significantly impacts how inspectees choose to behave.

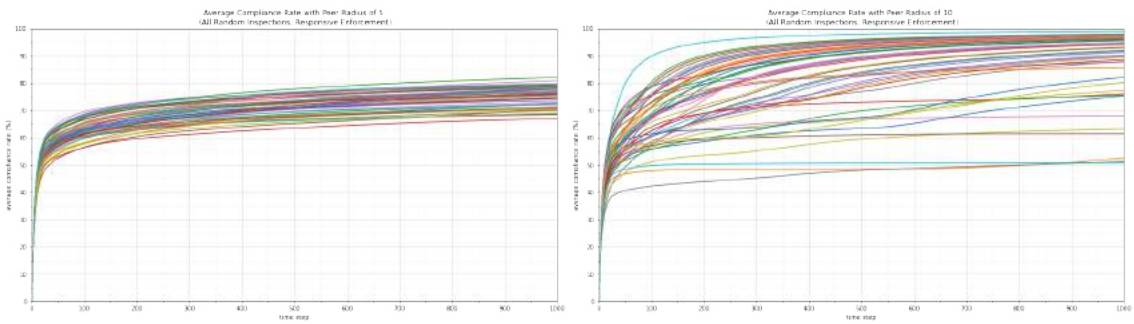

(a) r_{peers}=5.
For a breakdown by compliance level, see Figure A.70.

(b) r_{peers}=10.
For a breakdown by compliance level, see Figure A.71.

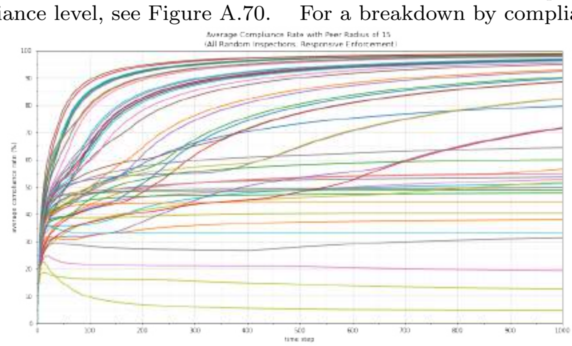

(c) r_{peers}=15.
For a breakdown by compliance level, see Figure A.72.

Figure 5.42: Average compliance rate results (r_{peers}: 5, 10, & 15; RE).

Table 5.3 summarizes the results of the sensitivity analysis and structural validation.

	Structurally Valid Ranges	
Parameter	Standard Enforcement	Responsive Enforcement
Initial share of compliant inspectees ($\%_{init,comp}$) *AR inspection strategy & r_{peers}=10*	42.5% - 47.5%	38% - 40%
Percent of Random Inspections ($\%_{rand}$) *MRRbNC inspection strategy*	10% - 90%[a]	10% - 90%
Percent of Random Inspections ($\%_{rand}$) *MRRbNC inspection strategy*	10% - 90%[a]	10% - 90%
Logit Normal Parameters (μ and σ) *AR inspection strategy & r_{peers}=10*	See Table 5.2	
Radius of Peers (r_{peers}) *AR inspection strategy*	5[b]	5 - 15

Table 5.3: Structurally valid ranges of varied parameters in the sensitivity analysis.
[a] This range of $\%_{rand}$ produces an inspectee population with more than 50% of non-compliant inspectees. Though this is not in alignment with definition of structural validity as described in the beginning of Section 5.1, the model still behaves as expected and does not converge unrealistically towards either zero or 100% compliance.
[b] Similarly, a radius of five with SE produces an inspectee population with more than 50% of non-compliant inspectees. However, this is the only viable value for the radius, as increasing it further only decreases the compliance rate.

5.2 PABMI: Scenarios

The aim of the scenario analysis is to find the most effective intervention strategies given different behavioral characteristics of the inspectee population. With unknown aspects of the inspection environment, these scenarios spur the ILT to explore and visualize possible future dynamics given assumptions about inspectees' behavior and behavioral trajectories found in data. By depicting unknown aspects of the inspection environment, the PABMI can be used to help policymakers develop interventions that are effective on most occasions.

The scenarios are summarized in Table 5.4. For each scenario, all five types of inspection strategy (see Table 3.4) are simulated and compared to one another to see which produces the best outcomes of interests. Each inspection strategy in each scenario was simulated 50 times. The share of compliant inspectees is used as a proxy of good behavior in the inspectee population, and the time it takes for the inspectee population to become compliant is considered.

Scenario	React to Peer Pressure?	React to Inspections?	React to Enforcement?
SIM-1: Individualistic, Non-responsive	No	Yes	No
SIM-2: Networked, Non-responsive	Yes	Yes	No
SIM-3: Individualistic, Responsive	No	Yes	Yes
SIM-4: Networked, Responsive	Yes	Yes	Yes

Table 5.4: PABMI scenarios summary.

The scenarios test two dimensions of real-life factors that impact inspectees' behavior: peer pressure effects and reaction to enforcement. These two dimensions are behavioral phenomena that are observed by inspectors but lack the adequate hard data to conduct statistical analysis for a purely phenomenological approach (see Section 2.1.2). Therefore, these scenarios were crafted to see how the model behaves with different combinations of behavioral phenomena. An individualistic inspectee population reflects an environment where behavior is not affected by peer pressure. Conversely, a networked population contains inspectees whose behavior is influenced by their peers. A non-responsive inspectee population means that inspectees are not influenced by the enforcement intervention, while a responsive population reacts to the enforcement intervention when it is doled out. In all the scenarios, only RE is simulated, as the sensitivity analysis shows that it consistently yields better outcomes than SE. The reaction to inspections is always switched on, as this is the only behavioral dynamic that was calibrated by statistical patterns found in the Inspectieview Binnenvaart data. Table 5.5 shows the parameter values used for the scenarios. The percent of random inspections, logit normal mu, logit normal sigma, and radius of peers were determined based on the results of the sensitivity analysis. The remaining parameters were determined phenomenologically from empirical data.

Parameter	Value	Scenario
Number of inspectees	7878	All
Length of run	1000	All
% of inspectees who are compliant	39.6%	All
% of the non-compliant population who are unintentional offenders	20.8%	All
% of the non-compliant population who are conscious offenders	76.4%	All
% of the non-compliant population who are criminal offenders	2.8%	All
Average number of violations for unintentional offenders	3.99	All
Average number of violations for conscious offenders	14.67	All
Average number of violations for criminal offenders	0.54	All
Number of inspectors	100	All
Number of inspections per inspector	13.13	All
% of random inspections[a]	50%[a]	All[a]
% of the non-compliant population who are escalatory	2.97%	All
% of the non-compliant population who are deescalatory	6.98%	All
Logit normal mu (μ) parameter for absorbance capacity	1.0[b]	SIM-3 & 4
Logit normal sigma (σ) parameter for absorbance capacity	0.37[b]	SIM-3 & 4
Radius of peers	10	SIM-2 & 4

[a] Only for mixed inspection strategies (see Inspection Strategies #4 and #5 in Table 3.4).
[b] See Table 5.2.

Table 5.5: PABMI model parameter values for scenarios.

5.2.1 SIM-1: Individualistic, Non-responsive

In this first scenario, inspectees do not react to peer pressure or enforcement interventions; they only react to inspections. This scenario reflects the most phenomenological state of the PABMI, as none of the parameters in this scenario are based on behavioral theories. In this case, the RbNC inspection strategy is the most effective in the short term (Figure 5.43). In the long haul, the MRRbOS inspection strategy is the most effective at increasing the share of compliant inspectees (Figure 5.43). However, the MRRbOS inspection strategy requires a long time horizon, approximately 125 inspection cycles (Figure 5.43). Assuming two inspections are conducted per year, it would take 62.5 years before the MRRbOS inspection strategy produces the same share of compliant inspectees as the RbNC inspection strategy (Figure 5.43).

The RbNC inspection strategy is effective because in the inland shipping sector, the worst compliance record is held by conscious violators who make up three-quarters of the non-compliant population (see Table 5.5). Thus, targeting them will improve the share of compliant inspectees quickly because there is a large number of them. However, the outcomes of this strategy quickly plateau as original unintentional and criminal violators are targeted less often, if at all, with this strategy. This limits the improvement of the compliance rate over time. The MRRbOS inspection strategy takes a longer time horizon than the RbNC inspection strategy because deescalatory criminal offenders require three inspection cycles to become compliant while deescalatory conscious offenders only require two inspection cycles. Therefore, targeting risk-based inspections based on offense severity (mainly criminal offenders) takes a longer time for the effect of inspections to be seen. Once the criminal offenders become lower-level offenders and eventually compliant, there is a phase shift where the population becomes increasingly more compliant.

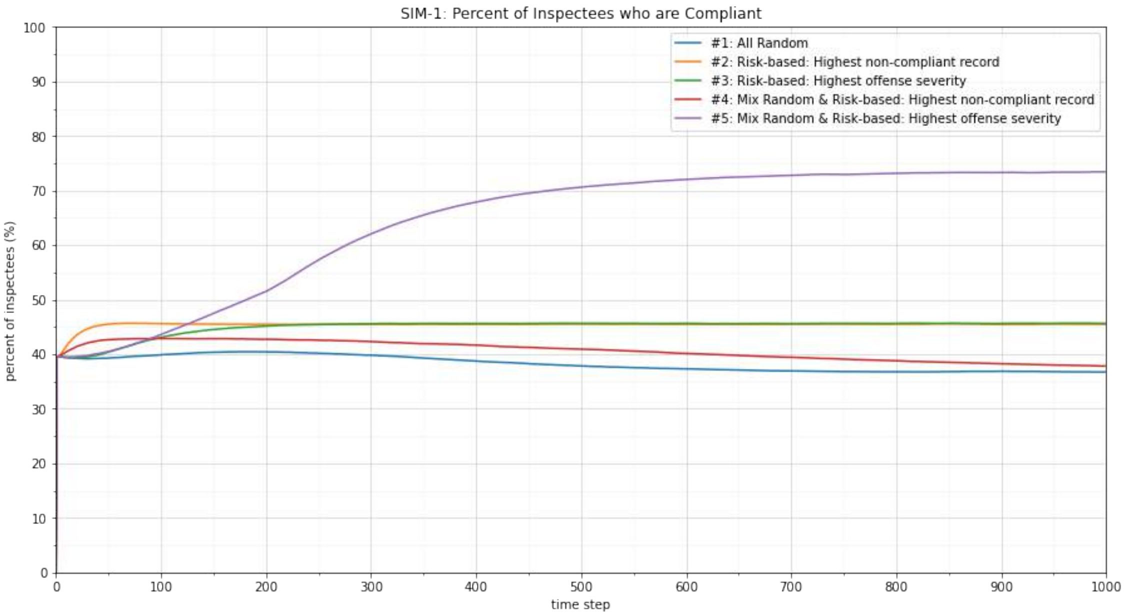

(a) 1000 ticks

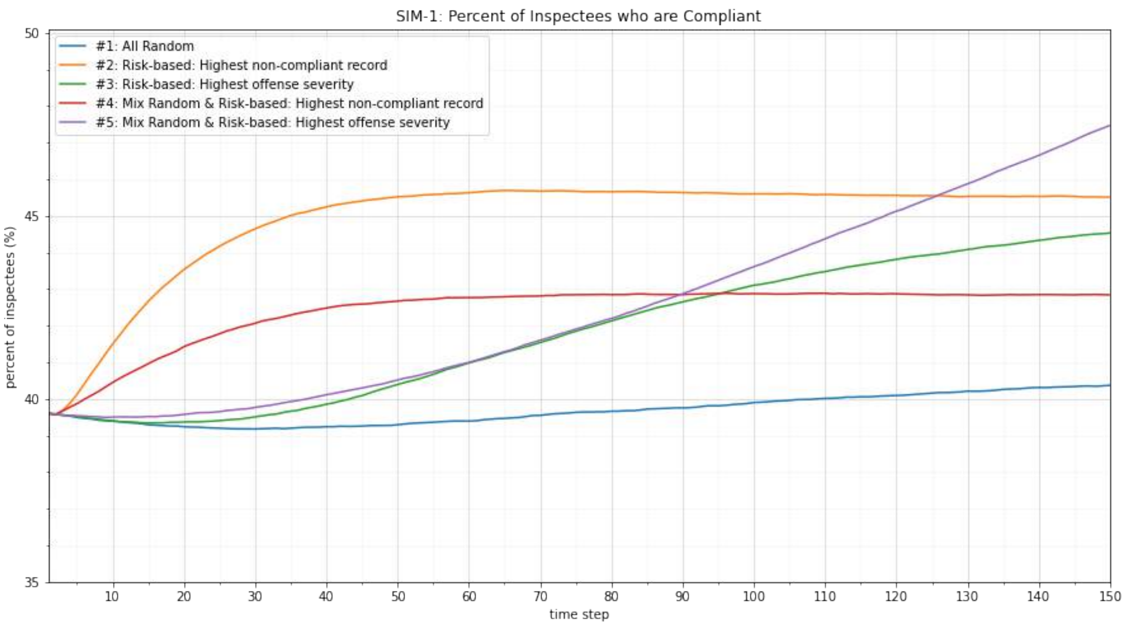

(b) 150 ticks

Figure 5.43: SIM-1: Share of compliant inspectees over time.

The average compliance rate of the population plateaus quickly for all the inspection strategies (Figure 5.44). The RbNC inspection strategy yields the worst average compliance rate (Figure 5.44). Because of the big difference between the historical number of violations between conscious offenders and unintentional and criminal offenders in the inland shipping sector (see Table 5.5), targeting inspections based on non-compliant record likely limits inspections to only conscious violators. This constrains the improvement of the average compliance rate of the inspectees when only one type of violator is targeted. The other inspection strategies allows for more inspections of inspectees with different violation types, whether by random inspections or targeting based on offense severity; this allows a wider range of inspectees to be inspected and therefore, a greater improvement in the average compliance rate.

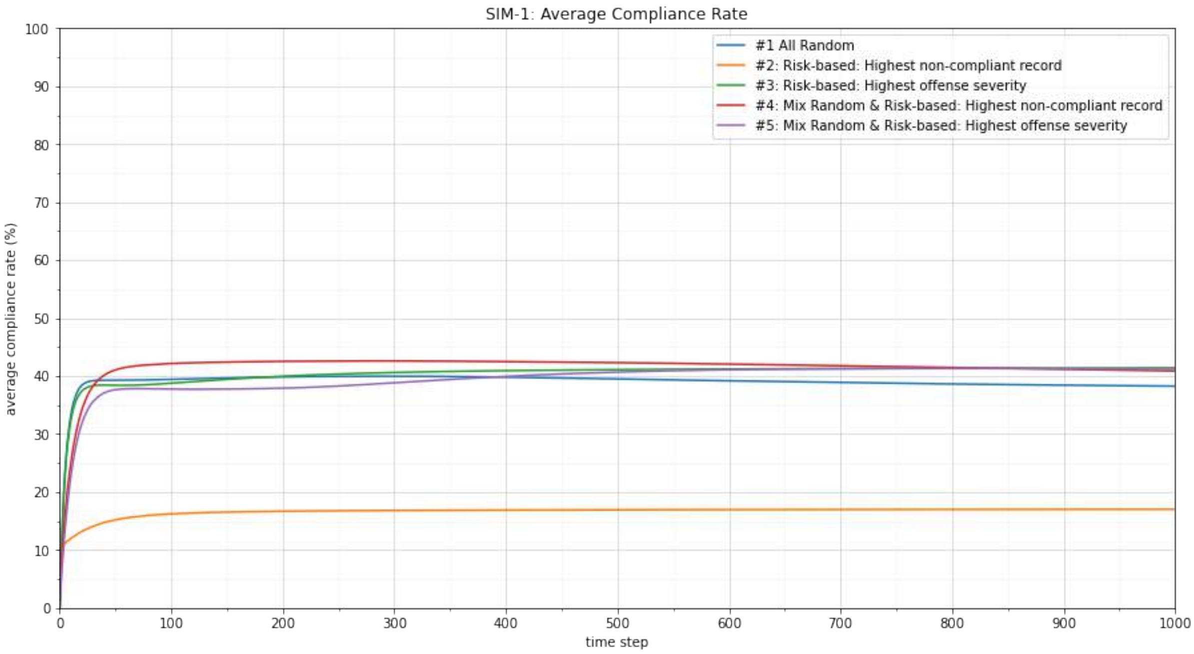

Figure 5.44: SIM-1: Average compliance rate results.

For all inspection strategies, there is an increase in the share of criminal inspectees, showing the effects of escalatory behavior over time (Figure 5.45). Yet, this is only a small share of the non-compliant population. Deescalatory behavior is found to be more prevalent in the inland shipping sector (see Section 2.1.2), which is shown by the increase in unintentional violators as the number of conscious violators decreases (Figure 5.45). Notably, only the MRRbOS inspection strategy is capable of decreasing the share of unintentional violators even after the initial deescalation from conscious violating behavior (Figure 5.45a). This shows that by randomizing a portion of inspections, there is further potential for additional inspectees to deescalate their behavior until they become compliant. However, this requires a long time horizon. Still, the effect of inspections is small as shown by the plateau of the share of compliant inspectees at around 45.5% for the RbNC inspection strategy (Figure 5.43).

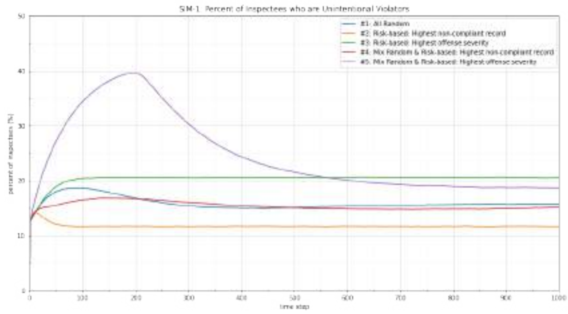

(a) Unintentional violators.

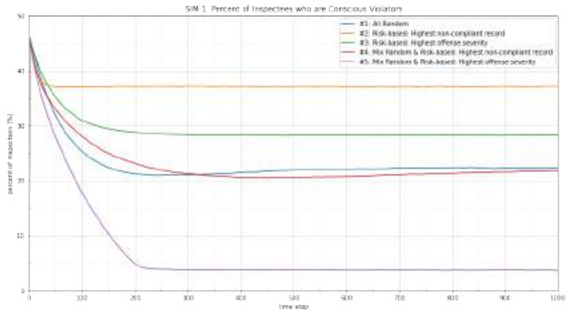

(b) Conscious violators.

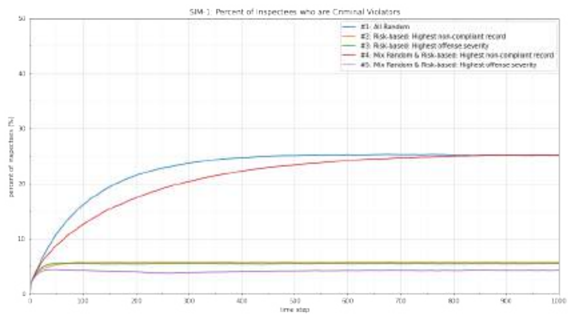

(c) Criminal violators.

Figure 5.45: SIM-1: Share of non-compliant inspectees over time.

Future research should test effectiveness of an inspection strategy that begins as a pure RbNC strategy and over time, starts to introduce a portion of risk-based inspections based on offense severity. To illustrate, this inspection strategy would be purely RbNC for 30 inspection cycles. Then, between 30 and 40 inspection cycles, 10% of inspections would be and 90% RbNC. Between 40 and 50 inspection cycles, 20% of inspections would be and 80% RbNC, and so on. The hypothesis is that starting off with a pure RbNC strategy would take advantage of the fast identification of conscious violators. As the effectiveness of the RbNC starts to plateau, introducing a portion of RbOS inspections would take advantage of picking off higher-level offenders to boost the share of compliant inspectees beyond the plateau.

5.2.2 SIM-2: Networked, Non-responsive

This scenario simulates a sector where inspectees are influenced by peer pressure and respond to inspections, but do not react to responsive enforcement. In this scenario, the most effective inspection strategy is the RbNC inspection strategy, followed by the MRRbNC inspection strategy (Figures 5.46 and 5.47). Targeting inspections based on non-compliant record is more effective in a networked environment compared to an individualistic one. Yet, in an networked environment, the inspection strategies alone were not enough to bring the inspectee population to a majority compliant position. These results show that given the majority initial share of non-compliant inspectees and the relatively smaller percentage of inspectees that exhibit deescalatory behavior (see Table 4.1), negative peer pressure overwhelms the effect of inspections. In other words, inspections are an inadequate intervention for increasing compliance behavior in a networked environment, as the effect of peer pressure by the majority non-compliant population is too strong.

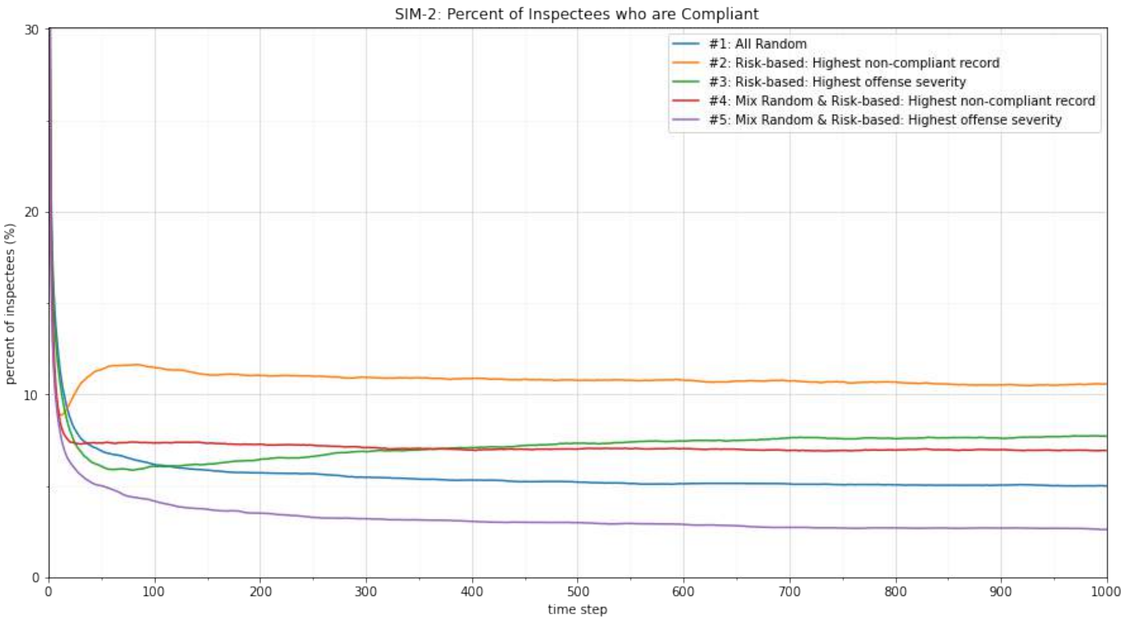

(a) 1000 ticks

Figure 5.46: SIM-2: Share of compliant inspectees over time.

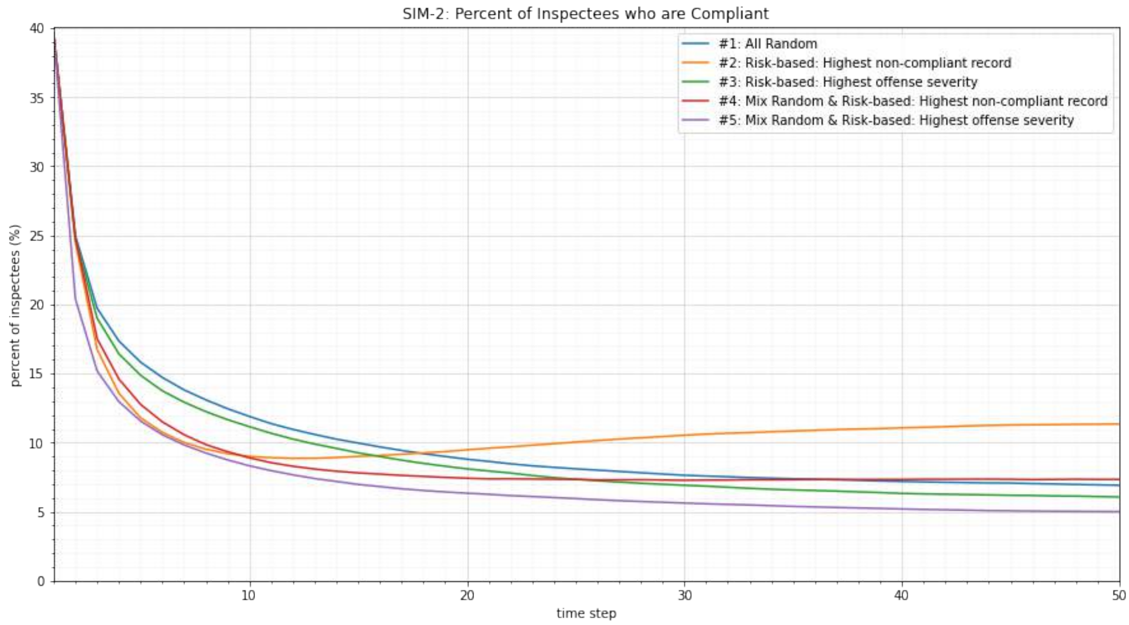

(b) 50 ticks

Figure 5.46: SIM-2: Share of compliant inspectees over time (continued).

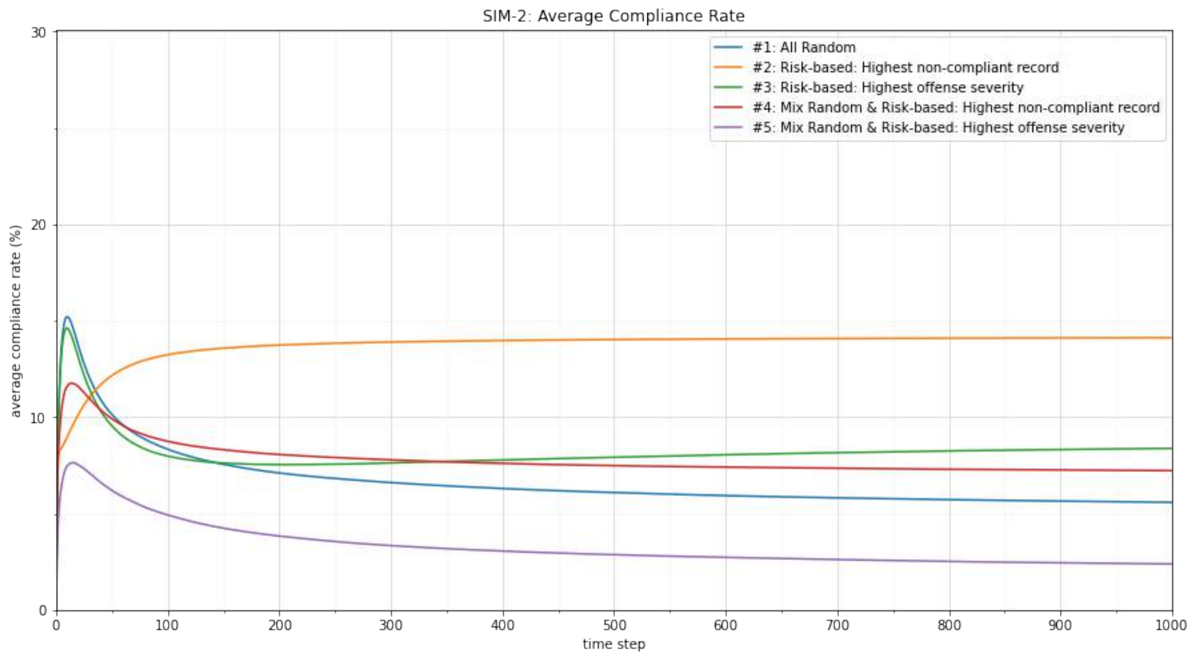

Figure 5.47: SIM-2: Average compliance rate results.

The difference between the share of conscious violators and unintentional or criminal violators widens from its initial values when the simulation begins (Figure 5.48). Because the initial population of non-compliant inspectees is majority conscious violators, a networked population will see a quick spread of conscious-level non-compliant behavior. The effect of inspections is not enough to counteract this peer pressure, leading to a majority non-compliant population.

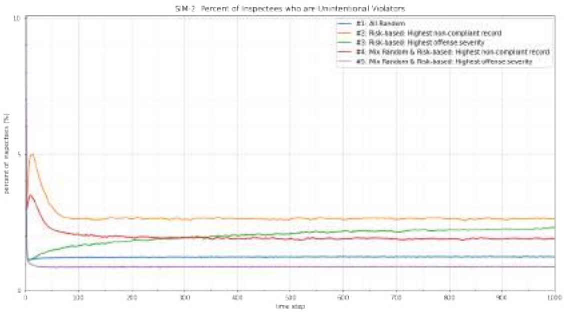

(a) Unintentional violators.

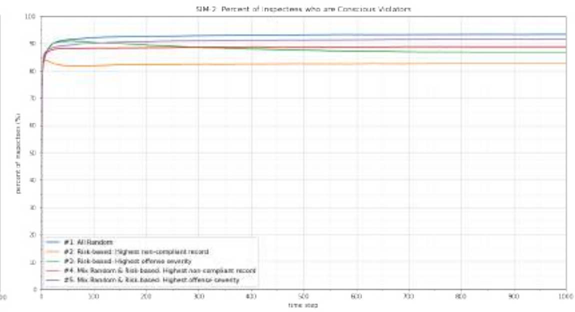

(b) Conscious violators.

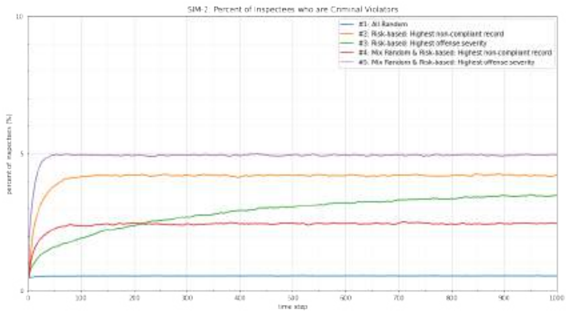

(c) Criminal violators.

Figure 5.48: SIM-2: Share of non-compliant inspectees over time.

The PABMI conceptualizes peer pressure as an inspectee conforming the most-occurring level of offense in its radius of peers (see Section 3.4.3). It assumes that any non-compliant inspectee that is inspected will be caught (see Section 4.1). In reality, it is possible that inspected non-compliant agents are not caught; the probability that a non-compliant inspectee gets caught and its effects on compliance outcomes are outside the scope of this research. Future research should include peers' learning behavior on how to avoid being caught during inspections. Escalatory and deescalatory behavior could be a symptom of an inspectee's calculation of their likelihood of being caught in the next inspection. In addition, other conceptualizations of the network using nodes and connectors with weights could be studied. However, since the population of non-compliant inspectees is larger than that of compliant inspectees, it is likely that another network conceptualization would still yield the same results.

5.2.3 SIM-3: Individualistic, Responsive

This scenario represents a sector where inspectees do not react to peer pressure; they only react to inspections and responsive enforcement. The most effective strategy for increasing the share of compliant inspectees is the MRRbOS inspection strategy in both the short and long term (Figure 5.49). The least effective strategy is the RbNC inspection strategy (Figures 5.49 and 5.50). These results show that targeting risk-based inspections based on offense severity is effective to reduce the non-compliance of higher level offenders first. As these offenders become compliant due to the effect of responsive enforcement, the risk-based inspections begin to target lower level offenders. The share of random inspections also allows for the consideration of a wide range of inspection candidates, which ensures that some lower level offenders are also inspected concurrently. This strategy casts a wide net while ensuring that the worst offenders are addressed early on. Compared to SIM-1, the rate at which the share of compliant inspectees increases is much higher. This is because criminal violators who cannot absorb the severity of enforcement become compliant in one inspection cycle; on the other hand, deescalatory criminal violators require three inspection cycles to become compliant. The effect of enforcement is much quicker than that of inspections.

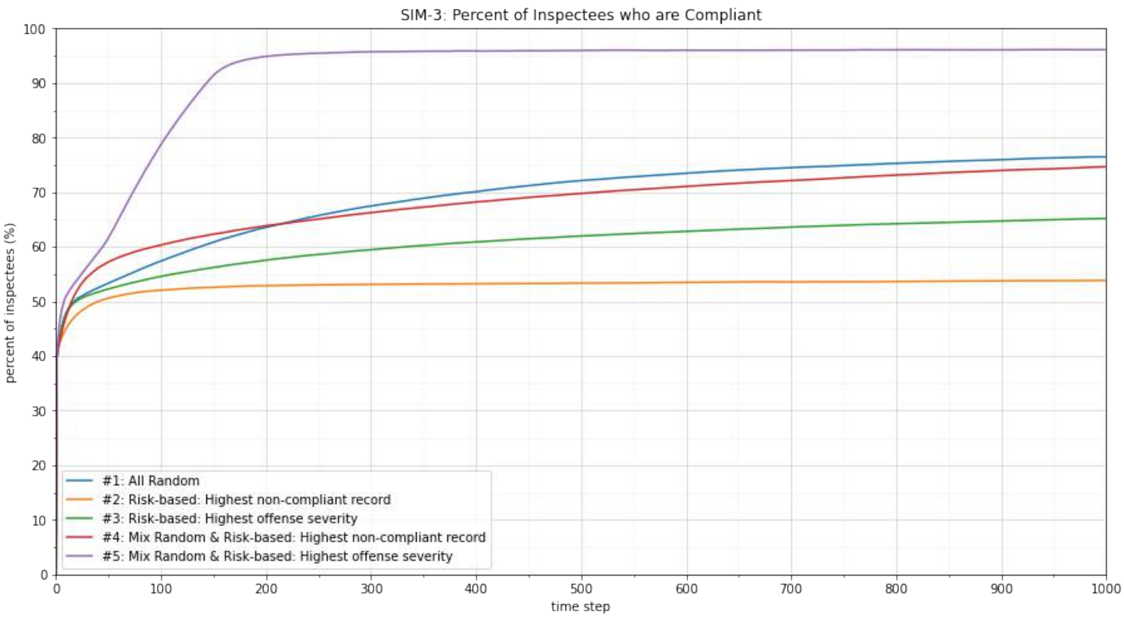

(a) 1000 ticks

Figure 5.49: SIM-3: Share of compliant inspectees over time.

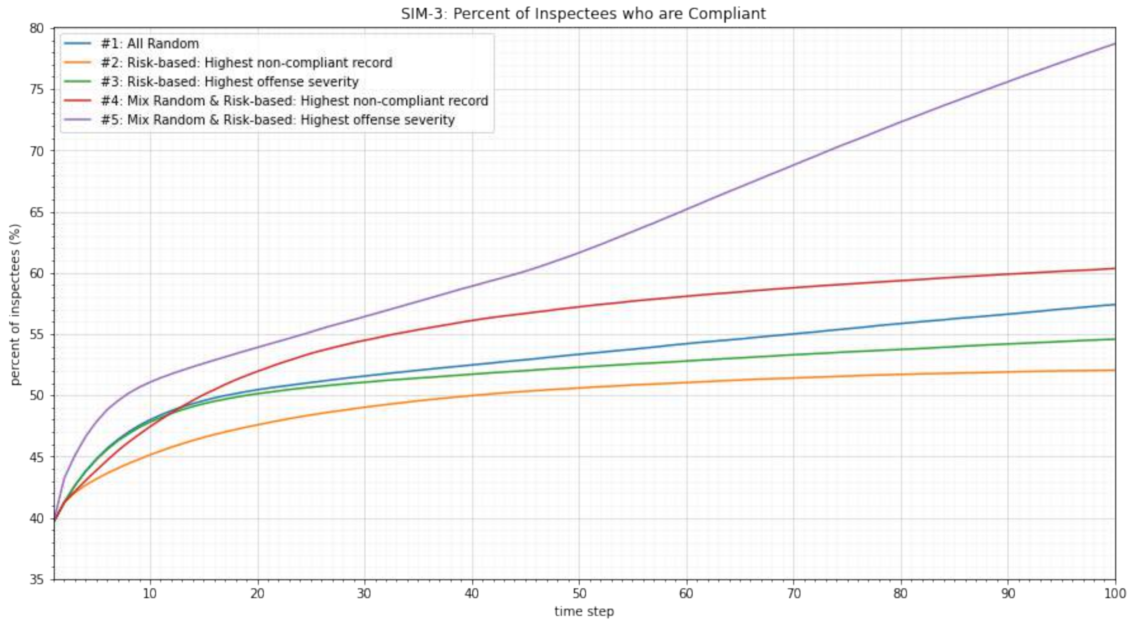

(b) 100 ticks

Figure 5.49: SIM-3: Share of compliant inspectees over time (continued).

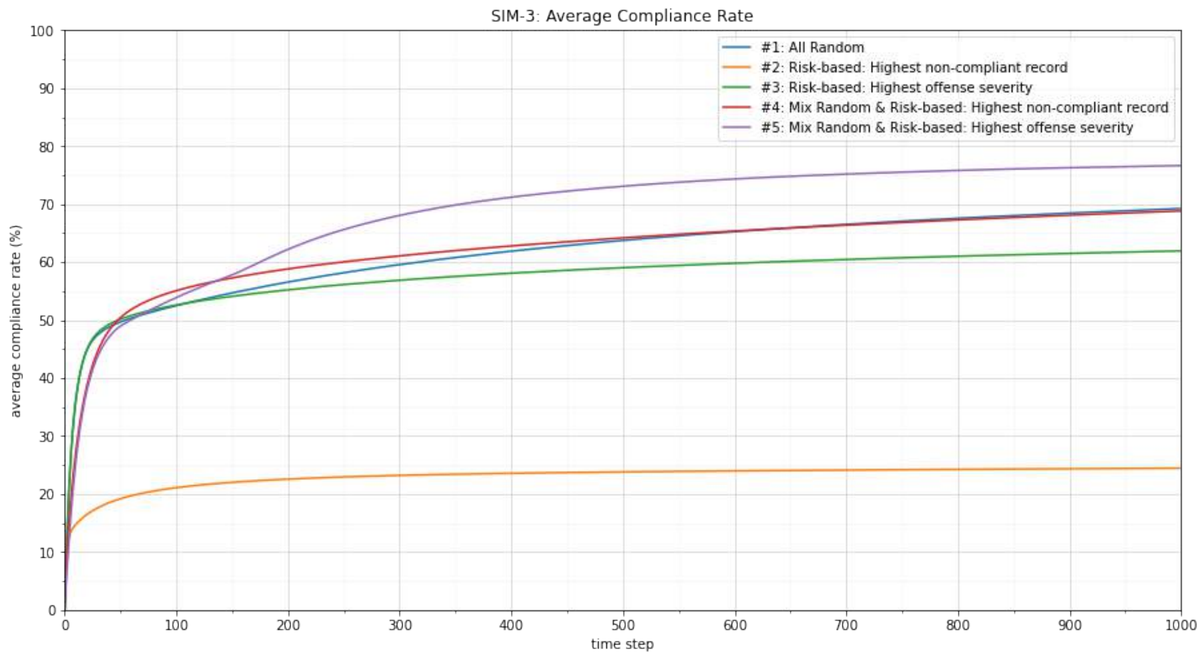

Figure 5.50: SIM-3: Average compliance rate results.

Because of the reaction to enforcement, the share of criminal inspectees is very low (Figure 5.51c). Similar to SIM-1, the deescalatory behavior shifts inspectees from being conscious violators (Figure 5.51b) to unintentional violators (Figure 5.51a), causing an increase in the share of unintentional violators at the start of the simulation. Notably, the MRRbOS inspection strategy is effective even in deescalating unintentional violators towards compliance, making it the most effective strategy in this scenario (Figure 5.51a). The reaction to enforcement allows the population of inspectees to move towards compliance quicker than inspections alone.

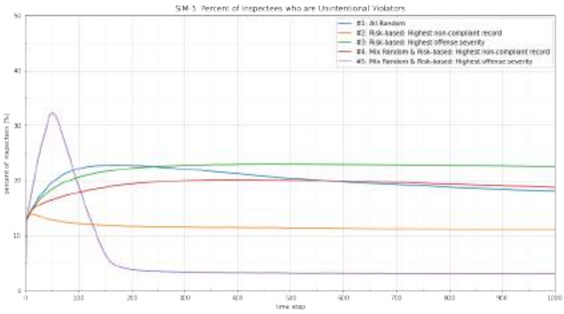

(a) Unintentional violators.

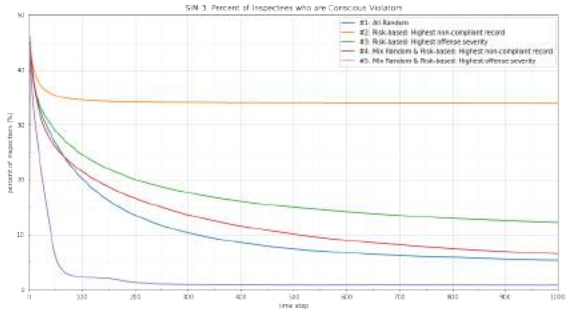

(b) Conscious violators.

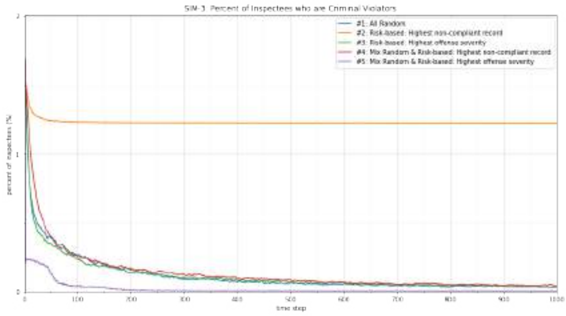

(c) Criminal violators.

Figure 5.51: SIM-3: Share of non-compliant inspectees over time.

5.2.4 SIM-4: Networked, Responsive

This scenario simulates a sector where inspectees react to peer pressure, inspections, and responsive enforcement. The most effective strategy for increasing the share of compliant inspectees is the MRRbOS inspection strategy (Figures 5.52 and 5.53). As mentioned in the preceding section, targeting inspection candidates based on offense severity is effective in responsive environments, as it addresses the highest level offenders quickly. In the long term, all inspection strategies except the AR inspection strategy yield at least a 95% share of compliant inspectees (Figure 5.52). However, the MRRbOS inspection strategy produces the fastest increase. Because this strategy contains a proportion of random inspections, there is a plateau of the share of compliant inspectees at 95%, while the RbOS strategy eventually yields a 100% compliant population. Though this is unrealistic, it shows that using a pure RbOS strategy systematically shifts the behavior of criminal, conscious, and unintentional violators in that order until all of them are compliant. Yet, this is redundant as it takes a long time horizon, and the MRRbOS inspection strategy already achieves more compliance in the near term.

Interestingly, both the RbNC and MRRbNC inspection strategies show a decrease in the compliant population before it increases after around 10 inspection cycles (Figure 5.52). The networked nature of the inspectees coupled with the high initial population of non-compliant inspectees leads to a sharing of non-compliant behavior that both inspections and enforcement cannot initially address. This "warm-up" period shows the effects of strong peer pressure at the start of the simulation. After some inspectees either deescalate their behavior after inspection or are unable to absorb enforcement, they become compliant. Then, there is a phase shift where these newly compliant inspectees can start to influence their peers in a positive, compliant way. This leads to the increase in the compliant share after the "warm-up" period of around 10 inspection cycles. Yet, even after the "warm-up" period, the RbNC and MRRbNC inspection strategies are not as effective as the MRRbOS inspection strategy.

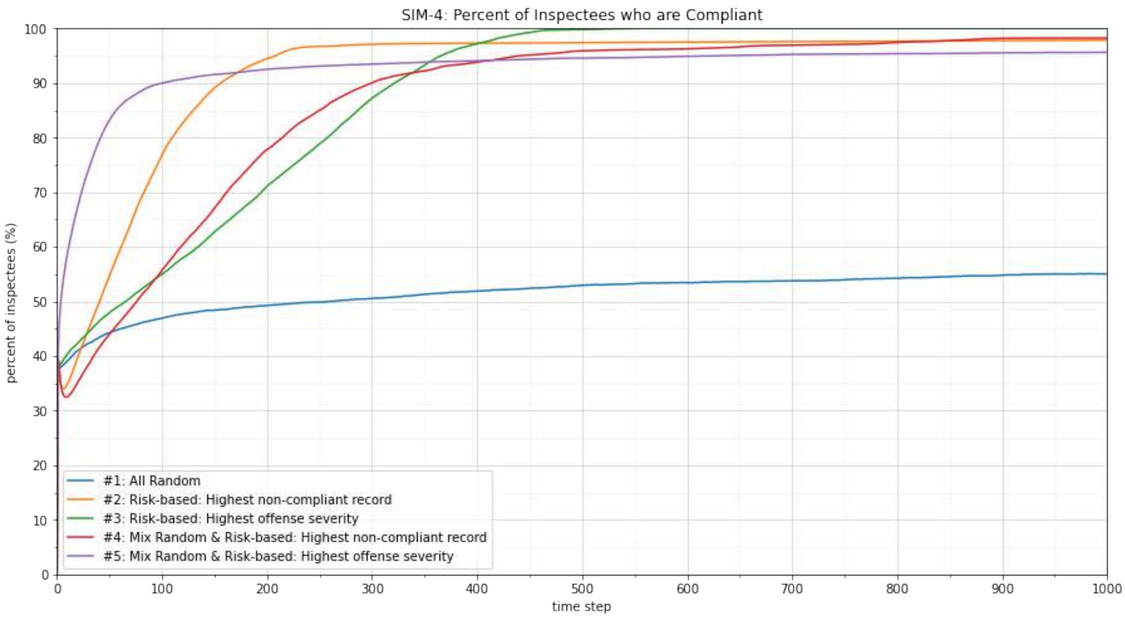

(a) 1000 ticks

Figure 5.52: SIM-4: Share of compliant inspectees over time.

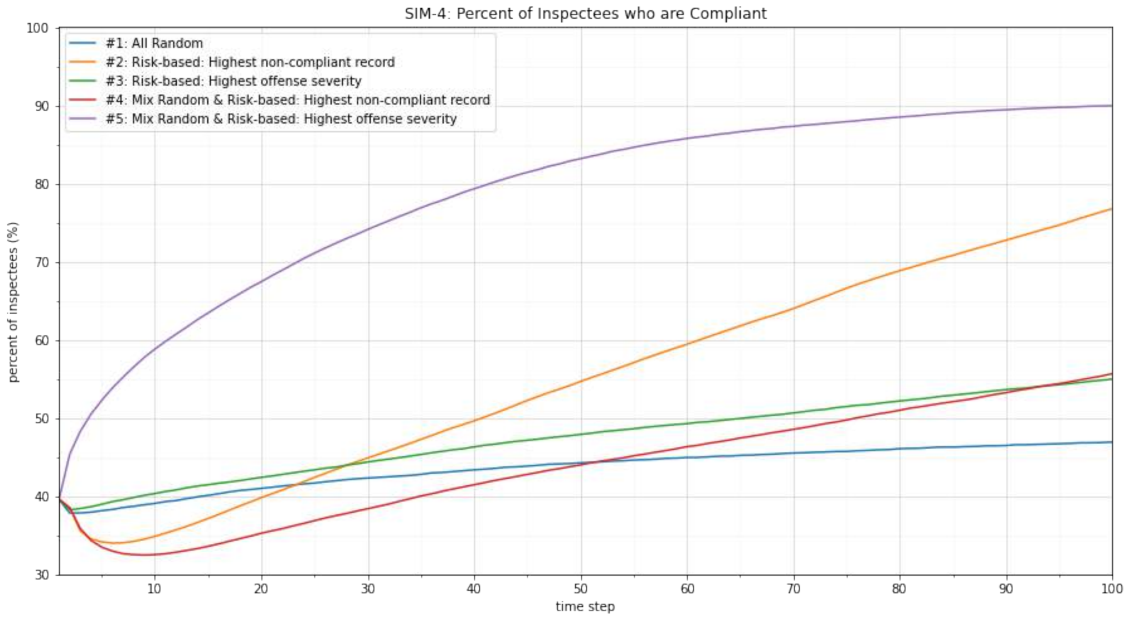

(b) 100 ticks

Figure 5.52: SIM-4: Share of compliant inspectees over time (continued).

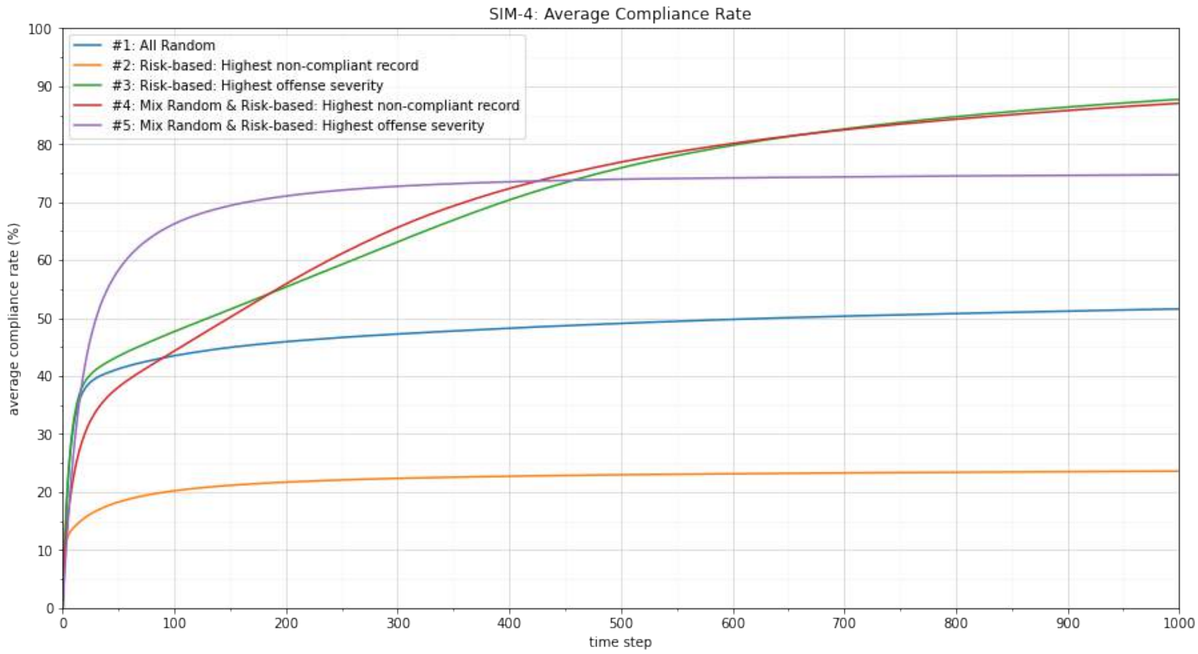

Figure 5.53: SIM-4: Average compliance rate results.

Because this scenario contains behavioral phenomena that is not based on empirical data, the simulation relies on the traditional approach to ABM to fill in the gaps in data, particularly on the effect of peer pressure and responsive enforcement (see Section 2.1.2). This blending of empirical behavioral phenomena and behavioral theories is reflected in the results, even though having a near 100% compliant population is unrealistic (Figures 5.52 and 5.54). These results indicate that the peer pressure effects and reaction to enforcement as conceptualized in the PABMI can overcome nearly all non-compliance in approximately 400 inspection cycles with all inspection strategies except the AR strategy. While this time horizon is large, it is implausible that compliance outcomes improve that drastically. The challenge of incorporating both theoretical and empirical elements into an ABM and verifying it with the real world remains. Nevertheless, these results show that in a networked environment, effective enforcement helps to spread compliant behavior quickly and is the only way to counteract negative peer pressure (compare results of SIM-4 to those of SIM-2).

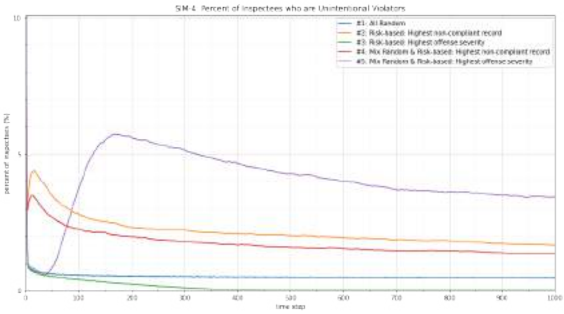

(a) Unintentional violators.

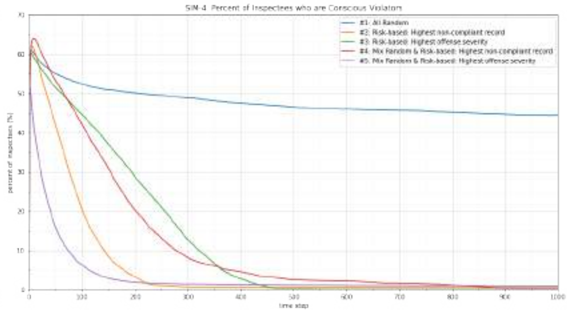

(b) Conscious violators.

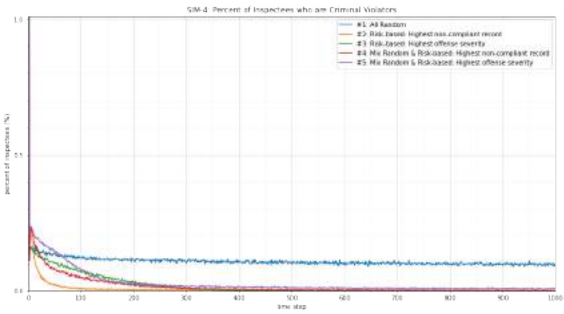

(c) Criminal violators.

Figure 5.54: SIM-4: Share of non-compliant inspectees over time.

5.2.5 Recommended Inspection Strategies

Table 5.6 shows the recommended inspection strategy by scenario.

Scenario	Recommended Inspection Strategy
SIM-1: Individualistic, Non-responsive	RbNC inspection strategy
SIM-2: Networked, Non-responsive	RbNC inspection strategy
SIM-3: Individualistic, Responsive	MRRbOS inspection strategy
SIM-4: Networked, Responsive	MRRbOS inspection strategy

Table 5.6: Summary of recommended inspection strategies by scenario.

These recommended strategies in Table 5.6 indicate the following:

1. For a non-responsive inspectee population, the RbNC inspection strategy is the most effective. This ensures that a large portion of non-compliant inspectees, particularly conscious offenders, are targeted at the outset. However, a networked population has negative peer pressure effects that are too strong, and it overcomes the effect of inspections. Therefore, the population of inspectees remains majority non-compliant. For an individualistic population, the MRRbOS is effective only after many inspection cycles. While it is useful to target criminal level offenders, it takes longer for higher-level offenders to deescalate their behavior until compliance. In this case, other mixed inspection strategies should be tested such as a mix of RbNC and RbOS strategies.
    - It is crucial to note that the PABMI was calibrated based on aggregated inland shipping data; this means that the historical non-compliant record for each type of offender was set to the average found in the empirical observations (see Sections 2.1.2 and 4.2). With an inspectee population with a large portion of conscious violators who also have the worst non-compliant record, targeting inspections based on historical non-compliant record covers a large portion of offenders. The effectiveness of the RbNC inspection strategy may be an artefact of the model conceptualization and is specific only to the inland shipping sector data. It is possible that the historical non-compliant record of inspectees does not follow a particular pattern, but because the empirical data was averaged, the inspectees were assigned a historical violations value based on a high level of aggregation.
2. For a responsive inspectee population, the MRRbOS inspection strategy is the most effective. Targeting criminal level offenders in a responsive environment means that many inspectees may not be able to absorb the high severity of

enforcement and therefore become compliant quickly. Meanwhile, maintaining a share of random inspections allows for a wider range of inspectees to be considered for inspection, increasing the chances that some may exhibit deescalatory behavior.

3. The effect of non-compliant peer pressure can only be counteracted by effective enforcement; inspections alone are insufficient for shifting the population towards compliance. While Table 5.6 shows that the best inspection strategy for SIM-2 is the RbNC inspection strategy, no inspection strategy can make the population majority compliant when there is peer pressure without a reaction to enforcement.

The results of the PABMI show inspectors which intervention strategies produce better outcomes, thereby helping them address their action dilemma by pinpointing where to focus their resources. First, the results of the PABMI confirm that inspections alone are not sufficient for increasing compliance rates. The ILT should seek to influence peer pressure in a positive way and enact enforcement interventions that nudge violators towards compliance, rather than relying solely on inspections. Inspectors should transition to either a RbNC or MRRbOS inspection strategy, instead of maintaining the status quo of inspecting offenders with highest historical non-compliance or only inspectees that they are familiar with. Repeat inspections of ships that were found to be non-compliant in previous inspections not only increases bias the Inspectieview Binnenvaart dataset, but also limits the pool of inspection candidates and hinders the effect of inspections and enforcement. Furthermore, inspectors should commit to a responsive enforcement strategy like the LHS, as standard enforcement based on principles of the "Interventieladder" (see Section 1.2.1.1) is not as effective. Yet, inspectors should monitor and start to collect data on reactions to responsive enforcement so that they can systematically test various enforcement interventions to find what is most effective for each inspectee subgroup. For the inland shipping sector, it is more worthwhile to devote resources to the exploration and data collection of peer pressure and enforcement influences, rather than the exploitation of resources on improving status quo operations. Verifying behavioral concepts such as peer pressure and responsive regulation is crucial for testing the effectiveness of the recommended policies that resulted from the PABMI analysis and for ILT's long-term inspection strategy.

Chapter 6

Discussion

Implementing a Phenomenological Approach to ABM

From a methodological standpoint, the main difference between a traditional and phenomenological approach to ABM is the identification of behavioral phenomena at the beginning of the research (see Table 1.4). A phenomenological approach requires quantitative data, quantitative data, or a combination of the two to be analyzed before the ABM is conceptualized. The behavioral patterns need to be identified in the data before it can be used to build the ABM. On the other hand, a traditional approach only requires a theory on the micro-interactions between agents, and the resulting behavioral patterns are found in the model outputs.

While the phenomenological approach to ABM aims to avoid relying on theories, it still exists adjacent to theories and incorporates them in an implicit way. For example, the depiction of an inspectee agent with state variables such as its location, non-compliant record, and likelihood of reacting to interventions are abstractions of reality; the assignment of state variables requires the modeler to make decisions on which characteristics are relevant for the study. These decisions are often formed based on behavioral theories on what factors are influential in shifting behavior (e.g., responsive enforcement; see Section 1.2.1). Even the most phenomenological state of the PABMI (see Section 5.2.1) incorporates inspectors' assumptions on what data points should be collected, which may have been informed by behavioral theories. To illustrate, inspectors gather data about an inland ship's inspection frequency, severity of the offense, and enforcement intervention. Implicit in this data collection method is the theory that inspection frequency, a violator's offense severity, and the enforcement severity are influential factors of compliance behavior; this is asserted by the theory of responsive regulation (see Section 1.2.1.1). In this sense, a phenomenological ABM is not completely free from theories.

Yet, a phenomenological ABM has its usefulness for inspectorates like the ILT. The statistical analysis of inspection data and its incorporation into the PABMI illuminate where there is still a lack of knowledge about the sector. As Epstein (2008) states, "without models...it is not always clear what data to collect." Yet, in the study of social systems, oftentimes theories inform data collection which can be analyzed later to confirm the existence of the theory (Epstein, 2008). This is the case for the development of the PABMI in this research. As illustrated in the preceding paragraph, the responsive regulation theory influenced the type of data that was collected, which were analyzed and incorporated to the PABMI. The results of the PABMI then further pinpoints areas that are still unknown and require further data collection, such as the peer pressure effects. Therefore, to make the most out of a phenomenological approach to ABM, the operations of an inspectorate should be an iterative process between theory, data collection, and modeling as shown in the Operational Framework for a Phenomenological Approach to ABM (Figure 6.1).

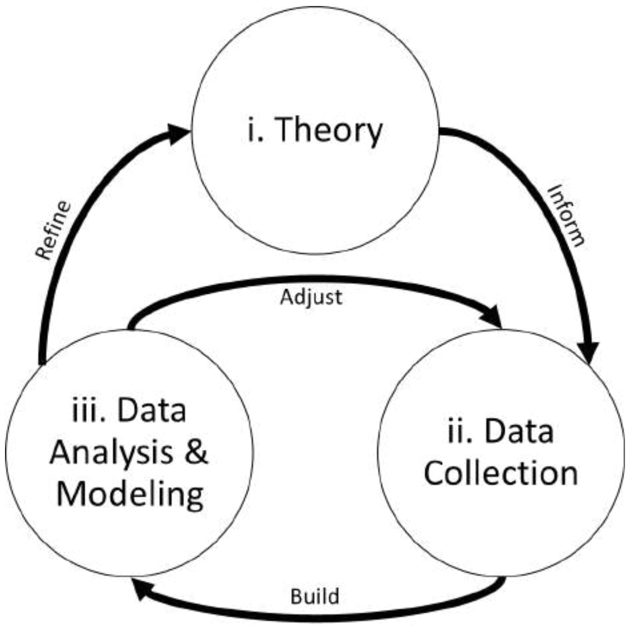

Figure 6.1: Operational Framework for a Phenomenological Approach to ABM: an iterative process between theory, data collection, and modeling.

Because inspectorates have bounded rationality (see Section 1.2.3.1), managing the iterative process between behavioral theories, data collection, and modeling is crucial for deploying resources efficiently; each one of these components plays a crucial role in the construction of the right type of information that inspectorates can use to address their action dilemma. To illustrate, behavioral scientists at the ILT can inform inspectors on what data to collect based on existing behavioral theories (*Inform* arrow). The qualitative and quantitative data collected by inspectors can be used by ILT's data scientists to build appropriate models (*Build* arrow). The results of the model should then refine existing behavioral theories (*Refine* arrow) and be communicated back to the inspectors so that they can adjust their data collection methods if necessary (*Adjust* arrow). This process allows inspectorates to address their action dilemma by 1) allocating their resources efficiently, 2) optimizing the right kind of data collection to fill any existing knowledge gaps, and 3) building models that test strategies and inform decisions on what interventions to deploy to increase compliance.

The scope of this research stops short of completing the full iterative cycle of the Operational Framework for a Phenomenological Approach to ABM (Figure 6.1), and ILT should dedicate resources to do so to make the most of the results of this research. This research only covers the blue portions of the Operational Framework as shown in Figure 6.2. Theories informed the type of data that was collected on the inland shipping sector (*Inform* arrow), and the collected inspection data was analyzed and used to build the PABMI (*Build* arrow). The results of the PABMI verify that the theory of responsive enforcement produces better compliance outcomes than standard enforcement; yet, it stops short of theorizing about the motivation behind why some inspectees react to inspection frequency (*Refine* arrow; see Section 5.2). The results of the PABMI must now be used to adjust existing data collection methods (*Adjust* arrow) to gather information on unknown aspects of the sector and validate the model assumptions. For example, the conceptualization of peer pressure and responsive enforcement should be validated. Once this data is collected, it be analyzed and incorporated into the model (*Build* arrow) so that it can again refine existing theories (*Refine* arrow). This process allows the ILT to more accurately investigate intervention strategies are the most effective for increasing compliance.

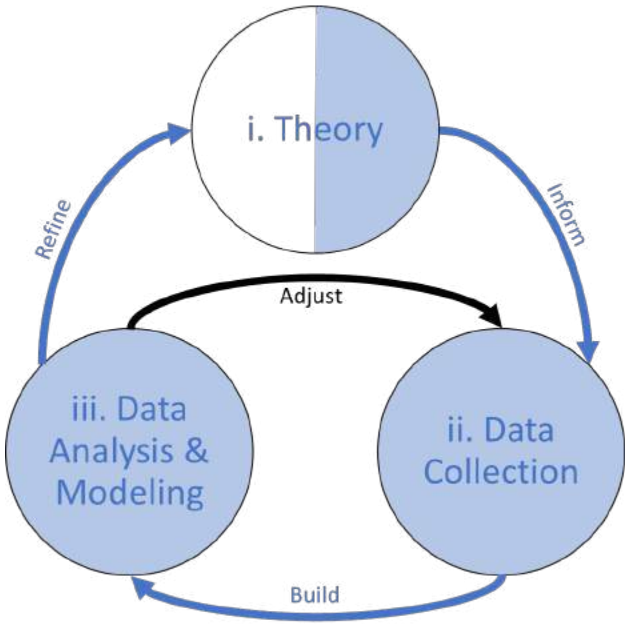

Figure 6.2: The Operational Framework for a Phenomenological Approach to ABM, where the blue arrows and shading represent what was completed in this research.

A phenomenological ABM itself (node iii in Figure 6.2) does not give insight into the theories of inspectees' behavioral motivations (node i in Figure 6.2); it does not generate possible micro-interactions that led to the observed, macro-level phenomena (see Table 1.4). There remains a need to discover the underlying intricacies of inspectees' behavior by linking the model results back to theory (*Refine* arrow in Figure 6.2). It is worth noting that the Operational Framework for a Phenomenological Approach to ABM shows that a pure phenomenological approach is not possible, as behavioral theories still play a role in the process. However, this research shows that the attempt at a phenomenological approach lessens the reliance on behavioral theories and helps to refine it efficiently over time. Additionally, it gives data a more integrated role in the development of an ABM. While traditional ABMs typically use empirical observations only at the end of the model development process (namely, to confirm the validity of the model outputs), a phenomenological ABM uses those empirical observations to build the model.

Risks of a Phenomenological Approach to ABM

The phenomenological approach to ABM generates assumptions that pose risks for the modeler and ILT practitioners. As mentioned in the preceding section, a phenomenological approach to ABM still incorporates some aspect of behavioral theories; for the PABMI the conceptualization of peer pressure effects and responsive enforcement relies on theories (see Section 2.1.2). In reality, the PABMI is a hybrid of a traditional and phenomenological approach to ABM (see Table 1.4). The risk of incorporating these theories is that they may not reflect the real-world, so they still need to be validated. For example, an agent's peers are determined by identifying neighboring agents within a radius of Euclidean space (see Section 3.4.3). Yet, there are other ways to characterize an inspectee's peers, such as a network and node model with weighted connectors. Moreover, the conceptualization of responsive enforcement relies on the theory that only inspectees that can absorb the severity of enforcement will continue to violate, while those that cannot will become compliant (see Sections 3.4.2 and 3.4.3). While a hybrid traditional and phenomenological ABM can fill in the gaps in data with the logic of ABM, further validation of the conceptualization of peer pressure and responsive enforcement is needed. To model responsive enforcement, the PABMI uses the *absorbance-capacity* state variable to characterize the inspectees' ability to withstand the severity of enforcement (see Section 3.2). For a phenomenological approach, this variable along with the radius of peers should be determined by data. This requires additional knowledge about inspectees' capacity to resist enforcement and the nature of the inland shipping network. To mitigate the risk of using components of behavioral theories, the modeler should liaise with ILT inspectors to conduct expert validation and inform future data collection efforts.

Furthermore, the data analysis process is not free from assumptions and biases in the interpretation of the results. For example, individual inspectors have discretion over how they assign violation types in the inland shipping dataset; there is no standard definition of what each type entails (see Section 2.1.2). Additionally, the GBTM analysis used to identify statistical behavioral phenomena averages the compliance characteristics of the inspectees over the entire time period that the data was collected (see Section 2.1.2). This high level of aggregation means that the model did not incorporate other distributions of agents' non-compliant record that could be present in the real world. For example, all unintentional inspectees were given a historical non-compliant record of 3.99 in the PABMI (see Table 2.7 and Section 3.2). However, unintentional inspectees likely vary in their historical non-compliant track record. The assumption that the historical non-compliant record of each inspectee subgroup can be characterized by an averaged value affects the results of the risk-based inspection strategies based on historical non-compliance (see Section 3.4.1). Moreover, there is a risk of overfitting the model to the data, given the calibration of the model to the highly-aggregated inspection data.

To make an ABM more phenomenological and integrative of empirical observations, the ILT would require more data; however, there is a risk of introducing additional bias into the inspection data (see Section 2.1.2.1). Moreover, the GBTM analysis that identifies behavioral trajectories before the model development process requires heavy time and resource commitments. The GBTM analysis only identifies one type of behavioral trajectory at a time, and there is no systematic way to test for which trajectories may exist. To mitigate the risk of wasting time and resources, the ILT should collaborate closely with inspectors to gather qualitative data on which variable would likely delineate behavioral subgroups of inspectees that exhibit different behavioral trajectories.

Applicability and Generalizability

Because of the calibration of the model to the Inspectieview Binnenvaart data, the PABMI is highly specific to the inland shipping sector. While the phenomenological parameters of the PABMI can be adjusted to recalibrate the model with the characteristics of other sectors, the aforementioned risk of overfitting the model to the inland shipping data makes repurposing the model for other sectors difficult. In addition, by nature of being based on observed inland shipping inspection data, it might not be applicable to other sectors (Van der Schaaf, 2019; see Section 1.2.2.4).

As the ILT does not have longitudinal inspection datasets of other sectors that are comparable to the Inspectieview Binnenvaart dataset, it is tough to determine whether the reaction to inspection found by the GBTM analysis is a behavioral phenomenon that also exists in other sectors. Additional data collection and behavioral trajectory analyses need to be conducted to ensure that inspectees of other sectors also react to inspections. If they do not, the PABMI cannot be generalized to that sector.

In addition, the PABMI assumes that the statistical relationship found between inspections and compliance behavior can be extended past the empirical observations of the Inspectieview Binnenvaart dataset to simulate possible futures. However, if the PABMI had been overfitted to the data, the relationship between inspections and compliance behavior cannot be accurately extrapolated beyond the duration of the Inspectieview Binnenvaart dataset. Nevertheless, the usefulness of the PABMI still stands. The PABMI reveals that the effect of inspections is smaller than the impact of peer pressure and enforcement (see Section 5.2.5). Given these impacts, the PABMI shows which strategies to pursue. Yet, the ILT should validate this with inspectors and pilot the recommended strategies to see whether it increases compliance in the real world (see Section 5.2.5). In the face of bounded rationality, the PABMI helps the ILT address their action dilemma as it integrates existing information at ILT's disposal and finds a "satisficing" strategy that optimizes the deployment of their resources.

Future Research Recommendations

Should inspectorates want to implement a phenomenological approach to ABM, they must first consider if they have sufficient data and adequate resources to conduct the upfront data analysis. Without identifying the macro-level behavior before the model development process, a phenomenological approach to ABM is not possible. For the ILT specifically, future research should focus on expanding the PABMI to include the learning behavior of inspectees. Learning behavior refers to the ways in which inspectees adjust their behavior to thwart inspections. The following methodological questions are recommended to frame future research.

1. What are the ways in which inspectees can thwart inspections?
    - To answer this question, future research should generate hypotheses on potential learning behavior of inland shipping inspectees. For example, inspectees could learn by observing the schedule of inspectors' visits and avoid being present during that time to evade inspections. They could

learn by observing the actions of their peers who successfully evade inspections. There could be informal or formal networks where inspectees share information about how to not get caught during inspections.

- To develop hypotheses, two methods can be used. First, a literature review can be conducted to identify how inspectees' learning behavior has been identified and studied in previous research. Second, unstructured and informational interviews to gather qualitative data from inspectors can bring to light any anecdotal evidence of learning behavior they observe in the sector.

2. What empirical data is needed to verify the existence of this learning behavior?
   - This question addresses the need for ILT to complete the *Adjust* arrow in the Operational Framework for a Phenomenological Approach to ABM (Figure 6.2) based on the results of the PABMI. In addition, this question also covers the *Inform* arrow in the framework (Figure 6.2), as the hypotheses generated in question #1 inform what data is needed.
   - First, inspectors must be engaged to validate the existing assumptions of the PABMI (*Adjust* arrow in Figure 6.2). Then, based on the hypotheses generated in question #1 and the data gaps identified from this research, inspectors can start to qualitatively and quantitatively collect those data points. For example, the nature of how inspectees are networked should be investigated further.
   - To collect data, inspectors can use industry-wide surveys to gain insight into how inspectees interact and share information. Additionally, inspectors can conduct interviews with inspectees to gather qualitative data on inspectees' perception of inspections.
3. How can the collected data be analyzed and conceptualized into an ABM?
   - To answer this question, the PABMI is improved and expanded to incorporate the outputs of question #2.
4. What inspection and enforcement strategies should be tested given the new expansion of the PABMI?
   - In addition to the inspection and enforcement strategies modeled in this research, other intervention strategies such as those proposed by Van der Heijden (2020) in Table 1.2 can be investigated by the expanded PABMI. This allows the researcher to evaluate which strategies are most effective at increasing compliance in an environment where inspectees exhibit learning behavior.

Chapter 7

Conclusion

This research evaluates the implementation and application of a phenomenological approach to ABM for determining intervention strategies that increase compliance in the Dutch inland shipping sector. The research sub-questions as shown below are addressed to answer the main research question: *How does a phenomenological approach to agent-based modeling avoid inaccurate presumptions on inspectees' behavioral motivations and show which inspection and enforcement strategies are most effective at increasing compliance among inspectees that exhibit idiosyncratic behaviors?*

Sub-question #1: What are the theoretical foundations for the application of phenomenological and agent-based modeling techniques for studying compliance behavior?

Previous studies use theories on the typologies of compliance to characterize patterns of causes behind an inspectees' particular behavior. These theories often assume that agents have perfect rationality where they always make decisions based on which option would optimize their desired outcome. However, perfect rationality is too simplistic, as it depicts inspectees as one-dimensional actors with access to perfect information and the capacity to analyze it. Moreover, it ignores the heterogeneity and agency of inspectees to act in counter-intuitive and idiosyncratic ways. Enforcement strategies such as ILT's Landelijke Handhavingsstrategie (LHS) aim to tailor interventions for each of the typologies of compliance, though the effectiveness of this responsive regulatory strategy still needs to be validated.

Empirical observations of behavior often contradict the typologies of compliance and may be noticed by inspectors before the underlying motivations are understood. Phenomenological modeling is a tool to investigate behavioral dynamics beyond existing typologies of compliance. Coupled with an agent-based approach that models agents with bounded rationality, phenomenological modeling provides a method of investing complexity starting from observation rather than theory. For the ILT, a phenomenological approach to ABM provides a tool for them to identify effective strategies that help them better address their action dilemma by pinpointing where and how to deploy their resources to achieve higher compliance rates.

Sub-question #2: What theories underpin the conceptualization of an ABM depicting the inspection environment?

Knol (2021) uses inspection game theory to build an ABM for the ILT. Her ABM contains two actors: inspectors and inspectees. In an inspection game, inspectors want to curb illegal activity of inspectees and detect violations, while inspectees choose whether to comply or not. Each actor is subject to time and resource constraints while they try to maximize their utility. Previous studies of inspection game theory show that it has limitations for modeling the inspection environment, mainly because inspectees in real life can be incentivized to act in ways that do not always maximize their utility.

Sub-question #3: What data is needed to identify observed behavioral phenomena in the inspectee population?

To circumvent the limitations of game theory and behavioral theories, longitudinal inspection data of the Dutch inland shipping sector and qualitative data by inspectors are needed to identify observed behavioral phenomena. The inspection data was analyzed using GBTM to find statistical evidence that inspectees' behavior are correlated with inspections. Anecdotal evidence by inspectors indicate that inspectees are also influenced by peer pressure and enforcement interventions.

Sub-question #4: How can behavioral phenomena be conceptualized into an ABM?

The behavioral phenomena identified in data acquired in sub-question #3 were used to define properties of inspectee agents in the PABMI. The model was calibrated with characteristics of the inland shipping sector, such as its size. In one tick that represents one inspection cycle, the ABM simulations three Go procedures where 1) inspectors choose who to inspect, 2) the chosen inspectees are inspected and the non-compliant inspectees are given an enforcement intervention, and 3) inspectees update their behavior.

In the first Go procedure, inspectors select inspection candidates by using any of the following five inspection strategies: 1) all random strategy, 2) risk-based using the highest non-compliant record, 3) risk-based using the highest offense severity, 4) mix of random and risk-based using the highest non-compliant record, and 5) mix of random and risk-based using the highest offense severity.

In the second Go procedure, inspectors inspect the selected candidates and record their level of compliance. The non-compliant inspectees are given an enforcement intervention. For standard enforcement, all non-compliant inspectees are given an enforcement intervention with a standard and consistent level of severity. For responsive enforcement, the severity of the enforcement intervention is commensurate the severity of the offense, as dictated by ILT's LHS.

In the third and final Go procedure, inspectees update their behavior based on three factors: peer pressure, inspection, and enforcement. Each of these factors can be switched on or off in the PABMI so that each one of these behaviors can be studied in isolation. To model peer pressure effects, inspectees conform to the most-occurring offense severity within their radius of peers. Inspectees' reaction to inspection is calibrated with the results of the GBTM analysis that found three main behavioral trajectories: 1) inspectees who become increasingly non-compliant with inspections, 2) inspectees who become increasingly compliant with inspections, and 3) inspectees who do not change their behavior with inspections. When inspectees react to enforcement, they assess whether they can withstand the severity of enforcement based on their absorbance capacity. If they can absorb the enforcement severity, they will continue to violate; if they cannot absorb it, they will become compliant.

Sub-question #5: How can the effectiveness of inspection and enforcement strategies be investigated under different scenarios of the inspectee population?

Four scenarios were simulated, reflecting a variation of two dimensions of real-life factors that impact inspectees' behavior: peer pressure effects and reaction to enforcement. These two factors were varied because they were discovered based on anecdotal evidence from inspectors, rather than statistical analysis of hard data. Meanwhile, the reaction to inspections is always switched on in the model because this behavior was discovered in hard data. The first scenario represents an inspectee population that only reacts to inspections and not peer pressure or enforcement (SIM-1: Individualistic, Non-responsive). This is the most phenomenological state of the PABMI as only empirical parameters are used. The second scenario depicts an inspectee population that reacts to peer pressure and inspections, but not enforcement (SIM-2: Networked, Non-responsive). In the third scenario, inspectees react to both inspections and enforcement, but not peer pressure (SIM-3: Individualistic, Responsive). In the fourth and final scenario, inspectees react to peer pressure, inspections, and enforcement (SIM-4: Networked, Responsive). Each of the five inspection strategies as detailed under sub-question #4 are simulated for each scenario along with responsive enforcement for SIM-2 and SIM-4.

Main research question: ***How does a phenomenological approach to ABM avoid inaccurate presumptions on inspectees' behavioral motivations and show which inspection and enforcement strategies are most effective at increasing compliance among inspectees that exhibit idiosyncratic behaviors?***

While this research set out to create a pure phenomenological ABM that avoids any reliance on behavioral theories, the PABMI turned out to be a hybrid traditional and phenomenological ABM that still utilizes components of behavioral theories. Inspectorates have bounded rationality, meaning they have incomplete knowledge about the inspection environment and a limited amount of resources to analyze available data. Where there was a lack of hard data, the modeler must rely on theories and the logic of traditional ABM. Yet, even the phenomenological aspects of the PABMI, namely the reaction to inspections, are not without its risks and limitations. Switching to a phenomenological approach simply trades the risks of the perfect rationality assumptions of behavioral theories with those of data analysis and interpretation. Though the PABMI avoids presumptions on the motivations of inspectees, it assumes that inspectees have behavioral mechanisms that are only supported by qualitative data. The ILT should validate the mechanisms of peer pressure and responsive enforcement in the PABMI. Nevertheless, the results of the PABMI provide the ILT with additional insight into the strategies that are most effective given incomplete information, enabling them to move closer to their desired outcome of increasing compliance while facing the action dilemma.

The results of the PABMI show that inspections alone are not sufficient for improving compliance behavior. The ILT should seek to influence peer pressure and provide enforcement interventions that change non-compliant behavior. Inspectors should transition to either a RbNC or MRRbOS inspection strategy for non-responsive and responsive environments, respectively. In addition, they should fully commit to a responsive enforcement strategy where the severity of the enforcement interventions is commensurate with the severity of the offense. Yet, the ILT should devote more resources to the exploration and data collection of peer pressure effects and responsive enforcement in the inland shipping sector. The effect of negative peer pressure, especially with the large share of non-compliant inspectees, causes non-compliant behavior to spread easily; only effective enforcement can counteract it. Further research should investigate the peer pressure effects and learning behavior of inspectees while the ILT pilots the recommended inspection strategies.

For sectors without comprehensive inspection data akin to the Inspectieview Binnenvaart dataset, it does not make sense to implement a phenomenological approach to ABM. Without data, the model would be no different than a traditional

ABM. Nevertheless, before replicating the methods used in this research, the ILT should first validate theoretical assumptions and enrich the data that is collected from the inland shipping sectors. Additional GBTM analyses on the Inspectieview Binnenvaart dataset can be conducted so that the behavioral trajectories based on other variables can be identified; these newly-discovered trajectories can be incorporated into the PABMI to simulate possible futures with more phenomenological parameters. However, this is time- and resource-consuming. The recommended path forward for the ILT is to complete the full iterative cycle of the Operational Framework for a Phenomenological Approach to ABM (Figure 6.2). This efficiently improves the model without sacrificing the realism provided by real data nor wasting time with upfront data analysis that may not be consequential.

This research contributes to existing scientific literature by applying two modeling techniques hardly used in the study of the inspection environment: phenomenological modeling and agent-based modeling. The methods used to combine data and theory to create the PABMI resulted in the development of an iterative operational framework that inspectorates can use to apply theories to their work without sacrificing the realism provided by data (Figure 6.2). Called the Operational Framework for a Phenomenological Approach to ABM, it streamlines the model improvement process within an inspectorate's resource and time constraints. The iterative cycle of the framework refines existing behavioral theories, data collection methods, and models. Over time, the model is improved, making it a powerful tool for assessing the effectiveness of an inspectorate's strategies (D'Orsogna and Perc, 2015). Additionally, the framework spurs discussion within inspectorates and provides a structure for dialogue about unforeseen drivers of behavior (Durlauf, 1999).

Most importantly, this research provides a societal contribution by directing the ILT towards how they can practically move closer to fulfilling their mission of safeguarding the sustainability of society and the environment. Inspectorates have limited resources and face an endless number of societal challenges that can be addressed (Black and Baldwin, 2010). Their action dilemma – having to decide on what action to take without knowing for certain whether their actions lead to better social outcomes – prevents them from making absolute optimal decisions with perfect rationality (De Bruijn et al., 2007). This research provides a framework that inspectorates can use to iteratively improve their strategies with limited resources. With better strategies and streamlined processes, malign actors can be identified quickly and prevented from acting in a way that might harm society.

References

Adhikari, A., & DeNero, J. (2017). Computational and inferential thinking: The foundations of data science.

Alm, J., Sanchez, I., De Juan, A., et al. (1995). Economic and non-economic factors in tax compliance. *KYKLOS-BERNE*, *48*, 3–3.

Alves, L. G., Ribeiro, H. V., Lenzi, E. K., & Mendes, R. S. (2013). Distance to the scaling law: A useful approach for unveiling relationships between crime and urban metrics. *PloS one*, *8*(8), e69580.

Alves, L. G., Ribeiro, H. V., & Mendes, R. S. (2013). Scaling laws in the dynamics of crime growth rate. *Physica A: Statistical Mechanics and its Applications*, *392*(11), 2672–2679.

Andrei, A. L., Comer, K., & Koehler, M. (2014). An agent-based model of network effects on tax compliance and evasion. *Journal of Economic Psychology*, *40*, 119–133.

Avenhaus, R., Von Stengle, B., & Zamir, S. (2002). Chapter 51: Inspection games. *Handbook of Game Theory with Economic Applications, ed. by R. Aumann, and S. Hart*, *3*.

Ayres, I., & Braithwaite, J. (1992). *Responsive regulation: Transcending the deregulation debate*. Oxford University Press, USA.

Baldwin, R., & Black, J. (2008). Really responsive regulation. *The modern law review*, *71*(1), 59–94.

Ball, P. (2012). *Why society is a complex matter: Meeting twenty-first century challenges with a new kind of science*. Springer Science & Business Media.

Barbaro, A. (2015). A place for agent-based models. Comment on "Statistical physics of crime: A review" by MR D'Orsogna and M. Perc. *Physics of Life Reviews*, *12*, 24–25.

Barra, A., Contucci, P., Sandell, R., & Vernia, C. (2014). An analysis of a large dataset on immigrant integration in Spain. The statistical mechanics perspective on social action. *Scientific reports*, *4*(1), 1–7.

Becker, G. S. (1968). Crime and punishment: An economic approach. Springer.

Becker, G. S., & Stigler, G. J. (1974). Law enforcement, malfeasance, and compensation of enforcers. *The Journal of Legal Studies*, *3*(1), 1–18.

Beinhocker, E. D. (2006). *The origin of wealth: Evolution, complexity, and the radical remaking of economics*. Harvard Business Press.

Bird, C., & Kortanek, K. (1974). Game theoretic approaches to some air pollution regulation problems. *Socio-Economic Planning Sciences*, *8*(3), 141–147.

Black, J., & Baldwin, R. (2010). Really responsive risk-based regulation. *Law & policy*, *32*(2), 181–213.

Borch, K. (1982). Insuring and auditing the auditor. *Games, economic dynamics, and time series analysis*, 117–126.

Braithwaite, J. (2016). *Responsive Regulation*. Retrieved April 12, 2022, from http://johnbraithwaite.com/responsive-regulation/

Brams, S., & Kilgour, D. M. (1988). *Game theory and national security*. Basil Blackwell, New York.

Brantingham, P. J., & Brantingham, P. L. (1984). *Patterns in crime*. Macmillan New York.

Brock, W. A., & Durlauf, S. N. (2001). Discrete choice with social interactions. *The Review of Economic Studies*, *68*(2), 235–260.

Burby, R. J., & Paterson, R. G. (1993). Improving compliance with state environmental regulations. *Journal of Policy Analysis and Management*, *12*(4), 753–772.

Castellano, C., Fortunato, S., & Loreto, V. (2009). Statistical physics of social dynamics. *Reviews of modern physics*, *81*(2), 591.

Chainey, S., & Ratcliffe, J. (2013). *GIS and crime mapping*. John Wiley & Sons.

Chalfin, A., & McCrary, J. (2017). Criminal deterrence: A review of the literature. *Journal of Economic Literature*, *55*(1), 5–48.

Christian, C. (2017). Enhanced enforcement outcomes through a responsive regulation approach to sales tax enforcement. *Journal of Public Budgeting, Accounting & Financial Management*.

Connolly, S. R., Keith, S. A., Colwell, R. K., & Rahbek, C. (2017). Process, mechanism, and modeling in macroecology. *Trends in Ecology & Evolution*, *32*(11), 835–844.

Contucci, P., & Ghirlanda, S. (2007). Modeling society with statistical mechanics: An application to cultural contact and immigration. *Quality & Quantity*, *41*(4), 569–578.

Contucci, P., & Vernia, C. (2020). On a Statistical Mechanics Approach to Some Problems of the Social Sciences. *Frontiers in Physics*, *8*, 585383.

De Bruijn, H., ten Heuvelhof, E., & Koopmans, M. (2007). *Law enforcement: The game between inspectors and inspectees*. Universal-Publishers.

D'Orsogna, M. R., & Perc, M. (2015). Statistical physics of crime: A review. *Physics of life reviews*, *12*, 1–21.

Durlauf, S. N. (1999). How can statistical mechanics contribute to social science? *Proceedings of the national academy of sciences*, *96*(19), 10582–10584.

Edmonds, B. (2017a). Different modelling purposes. *Simulating social complexity* (pp. 39–58). Springer.

Edmonds, B. (2017b). *Simulating social complexity*. Springer.

Edmonds, B., & Gershenson, C. (2015). Modelling complexity for policy: Opportunities and challenges. *Handbook on complexity and public policy*.

Epstein, J. M. (2008). Why model? *Journal of artificial societies and social simulation*, *11*(4), 12.

Étienne, J. (2010). Compliance theories. *Revue française de science politique*, *60*(3), 493–517.

Fortunato, S., Macy, M., & Redner, S. (2013). Editorial: Statistical mechanics and social sciences. *Stat. Phys*, *151*, 1–8.

Frigg, R., & Hartmann, S. (2020). Models in science (E. N. Zalta, Ed.). *The Stanford Encyclopedia of Philosophy (Spring 2020 Edition)*. https://plato-stanford-edu.tudelft.idm.oclc.org/archives/spr2020/entries/models-science/

Gräbner, C. (2016). Agent-based computational models–a formal heuristic for institutionalist pattern modelling? *Journal of Institutional Economics*, *12*(1), 241–261.

Grasmick, H. G., & Green, D. E. (1980). Legal punishment, social disapproval and internalization as inhibitors of illegal behavior. *J. Crim. L. & Criminology*, *71*, 325.

Gray, W. B., & Scholz, J. T. (1993). Does regulatory enforcement work-a panel analysis of OSHA enforcement. *Law & Society Review*, *27*, 177.

Güth, W., & Pethig, R. (1992). Illegal pollution and monitoring of unknown quality—a signaling game approach. *Conflicts and cooperation in managing environmental resources* (pp. 275–332). Springer.

Hayashi, C. (1998). What is data science? Fundamental concepts and a heuristic example. *Data science, classification, and related methods* (pp. 40–51). Springer.

Hegemann, R. A., Smith, L. M., Barbaro, A. B., Bertozzi, A. L., Reid, S. E., & Tita, G. E. (2011). Geographical influences of an emerging network of gang rivalries. *Physica A: Statistical Mechanics and its Applications*, *390*(21-22), 3894–3914.

Hilborn, R., & Mangel, M. (2013). *The Ecological Detective: Confronting Models with Data*. Princeton University Press.

Hoffer, L. D., Bobashev, G., & Morris, R. J. (2009). Researching a local heroin market as a complex adaptive system. *American journal of community psychology*, *44*(3), 273–286.

Hokamp, S., & Pickhardt, M. (2010). Income tax evasion in a society of heterogeneous agents–Evidence from an agent-based model. *International Economic Journal*, *24*(4), 541–553.

Höpfinger, E. (1979). Dynamic standard setting for carbon dioxide. *Applied game theory* (pp. 373–389). Springer.

Inspectie Leefomgeving en Transport. (n.d.). *Interventieladder*. Retrieved May 16, 2022, from https://web.archive.org/web/20190511080944/https://www.ilent.nl/over-ILT/handhaving-en-toezicht/interventie/interventieladder

Inspectie Leefomgeving en Transport. (2022). *Interventie*. Retrieved April 1, 2022, from https://www.ilent.nl/over-ilt/handhaving-en-toezicht/interventie

Islam, M. A., & McPhail, K. (2011). Regulating for corporate human rights abuses: The emergence of corporate reporting on the ILO's human rights standards within the global garment manufacturing and retail industry. *Critical Perspectives on Accounting*, *22*(8), 790–810.

Jones, P. A., Brantingham, P. J., & Chayes, L. R. (2010). Statistical models of criminal behavior: The effects of law enforcement actions. *Mathematical Models and Methods in Applied Sciences*, *20*(supp01), 1397–1423.

Kagan, R. A., & Scholz, J. T. (1980). The "criminology of the corporation" and regulatory enforcement strategies. *Organisation und recht* (pp. 352–377). Springer.

Keizer, K., Lindenberg, S., & Steg, L. (2008). The spreading of disorder. *science*, *322*(5908), 1681–1685.

Kendall, B. E., Briggs, C. J., Murdoch, W. W., Turchin, P., Ellner, S. P., McCauley, E., Nisbet, R. M., & Wood, S. N. (1999). Why do populations cycle? A synthesis of statistical and mechanistic modeling approaches. *Ecology*, *80*(6), 1789–1805.

Klages, A. (1968). *Spieltheorie und Wirtschaftsprüfung: Anwendung spieltheoretischer Modelle in der Wirtschaftsprüfung*. Appel.

Knol, A. (2021). Agent-based Model of an Inspection Game.

Korobow, A., Johnson, C., & Axtell, R. (2007). An agent–based model of tax compliance with social networks. *National Tax Journal*, *60*(3), 589–610.

Langbein, L., & Kerwin, C. M. (1985). Implementation, negotiation and compliance in environmental and safety regulation. *The Journal of Politics*, *47*(3), 854–880.

LaValle, S. M. (2006). 9.5.4 concerns Regarding Game Theory. *Planning algorithms.* Cambridge university press.

Llacer, T., Miguel, F. J., Noguera, J. A., & Tapia, E. (2013). An agent-based model of tax compliance: An application to the Spanish case. *Advances in Complex Systems*, *16*(04n05), 1350007.

Lodge, M., & Wegrich, K. (2012). *Managing regulation: Regulatory analysis, politics and policy.* Macmillan International Higher Education.

Macal, C. M., & North, M. J. (2005). Tutorial on agent-based modeling and simulation. *Proceedings of the Winter Simulation Conference, 2005.*, 14–pp.

Marin, S. R., Fernandez, B. P. M., & Vasconcelos, D. d. S. (2020). Agent Based Models and Complex Individuals: An Epistemological Analysis. *Œconomia. History, Methodology, Philosophy*, (10-2), 231–256.

May, P. J. (2005). Regulation and compliance motivations: Examining different approaches. *Public administration review*, *65*(1), 31–44.

McFadden, D. (2001). Economic choices. *American economic review*, *91*(3), 351–378.

Meester, J. (2021). *Development trajectories for Rule Violation in Dutch Inland Navigation: A Group-based Trajectory Analysis* (M. Eng. thesis). Leiden University.

Melchior, A., Dignum, F., & Ruiz, M. (2019). A Closer Look at Dutch Policy Development. *Conference of the European Social Simulation Association*, 383–395.

Ministerie van Justitie. (2006). De tafel van elf een veelzijdig instrument. *Retrieved from The Hague: www.minjus.nl.*

Mitchell, R. B. (2007). Compliance Theory, Compliance, Effectiveness, and Behaviour Change in International Environmental Law. *Oxford handbook of international environmental law*, 893–921.

Nagin, D. S. (2010). Group-based trajectory modeling: An overview. *Handbook of quantitative criminology*, 53–67.

Otto, S. P., & Day, T. (2011). *A Biologist's Guide to Mathematical Modeling in Ecology and Evolution.* Princeton University Press.

Paternoster, R., Saltzman, L. E., Waldo, G. P., & Chiricos, T. G. (1982). Perceived risk and social control: Do sanctions really deter? *LAw & Soc'y REv.*, *17*, 457.

Perc, M. (2019). The social physxics collective.

Picoli, S., Castillo-Mussot, M. d., Ribeiro, H. V., Lenzi, E., & Mendes, R. (2014). Universal bursty behaviour in human violent conflicts. *Scientific reports*, *4*(1), 1–3.

Rodrigue, N., & Philippe, H. (2010). Mechanistic revisions of phenomenological modeling strategies in molecular evolution. *Trends in Genetics*, *26*(6), 248–252.

Ruimschotel, D., Van Reenen, P., & Klaasen, H. (1996). De tafel van elf: Een conceptueel kader en een instrument bij rechtshandhavingsvraagstukken.

Schinckus, C. (2019). Agent-based modelling and economic complexity: A diversified perspective. *Journal of Asian Business and Economic Studies*.

Scholz, J. T. (1984). Cooperation, deterrence, and the ecology of regulatory enforcement. *Law and Society Review*, 179–224.

Scholz, J. T. (1997). Enforcement policy and corporate misconduct: The changing perspective of deterrence theory. *Law and Contemporary Problems*, *60*(3), 253–268.

Sherman, L. W. (1993). Defiance, deterrence, and irrelevance: A theory of the criminal sanction. *Journal of research in Crime and Delinquency*, *30*(4), 445–473.

Short, J. L., & Toffel, M. W. (2010). Making self-regulation more than merely symbolic: The critical role of the legal environment. *Administrative Science Quarterly*, *55*(3), 361–396.

Short, M. B., D'orsogna, M. R., Pasour, V. B., Tita, G. E., Brantingham, P. J., Bertozzi, A. L., & Chayes, L. B. (2008). A statistical model of criminal behavior. *Mathematical Models and Methods in Applied Sciences*, *18*(supp01), 1249–1267.

Simon, H. A. (1957). A behavioral model of rational choice. *Models of man, social and rational: Mathematical essays on rational human behavior in a social setting*, 241–260.

Smojver, S. (2012). Analysis of banking supervision via inspection game and agent-based modeling. *Central European Conference on Information and Intelligent Systems*, 355.

Stigler, G. J. (1970). The optimum enforcement of laws. *Journal of Political Economy*, *78*(3), 526–536.

Tsebelis, G. (1990). Penalty has no impact on crime: A game-theoretic analysis. *Rationality and Society*, *2*(3), 255–286.

Tsebelis, G. (1991). The effect of fines on regulated industries: Game theory vs. decision theory. *Journal of Theoretical Politics*, *3*(1), 81–101.

Tsebelis, G. (1993). Penalty and crime: Further theoretical considerations and empirical evidence. *Journal of Theoretical Politics*, *5*(3), 349–374.

Tyler, T. R. (1990). Why people obey the law. *Why people obey the law*. Yale University Press.

Van Asselt, E., Osinga, S., & Bremmers, H. (2016). Simulating compliance behaviour for effective inspection strategies using agent based modelling. *British Food Journal*.

Van Dam, K. H., Nikolic, I., & Lukszo, Z. (2012). *Agent-based modeling of socio-technical systems* (Vol. 9). Springer Science & Business Media.

Van der Heijden, J. (2020). Responsive regulation in practice: A review of the international academic literature. *State of the Art in Regulatory Governance Research Paper–2020.06*.

Van der Schaaf, A. (2019). Potentials and Limits of Phenomenological Models. *Modelling radiotherapy side effects: Practical applications for planning optimisation* (p. 52). CRC Press.

Van der Voort, H., Sedee, I., Booijink, T., & Vaart, E. (2020). Is datagedreven risicogebaseerd toezicht op termijn effectief? *Tijdschrift voor Toezicht*, *11*, 153–162. https://doi.org/10.5553/TvT/187987052020011003004

Van Duin, S. R., Dekker, H. C., Wielhouwer, J. L., & Mendoza, J. P. (2018). The tone from above: The effect of communicating a supportive regulatory strategy on reporting quality. *Journal of Accounting Research*, *56*(2), 467–519.

Van Erp, J. (2011). Naming without shaming: The publication of sanctions in the Dutch financial market. *Regulation & Governance*, *5*(3), 287–308.

Van Snellenberg, T., & Van de Peppel, R. (2002). Perspectives on compliance: Non-compliance with environmental licenses in the Netherlands. *European Environment*, *12*(3), 131–148.

Weissing, F., & Ostrom, E. (1991). Irrigation institutions and the games irrigators play: Rule enforcement without guards. *Game equilibrium models ii* (pp. 188–262). Springer.

White, C. R., & Marshall, D. J. (2019). Should we care if models are phenomenological or mechanistic? *Trends in Ecology & Evolution*, *34*(4), 276–278.

Willensky, U. (2021). *Netlogo*. https://ccl.northwestern.edu/netlogo/ Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL

Zhu, J., & Chertow, M. R. (2019). Authoritarian but responsive: Local regulation of industrial energy efficiency in Jiangsu, China. *Regulation & Governance*, *13*(3), 384–404.

Appendix A

Appendix

A.1 PABMI: Sensitivity Analysis Results

A.1.1 Varying Initial Share of Compliant Inspectees ($\%_{init,comp}$)

Standard Enforcement

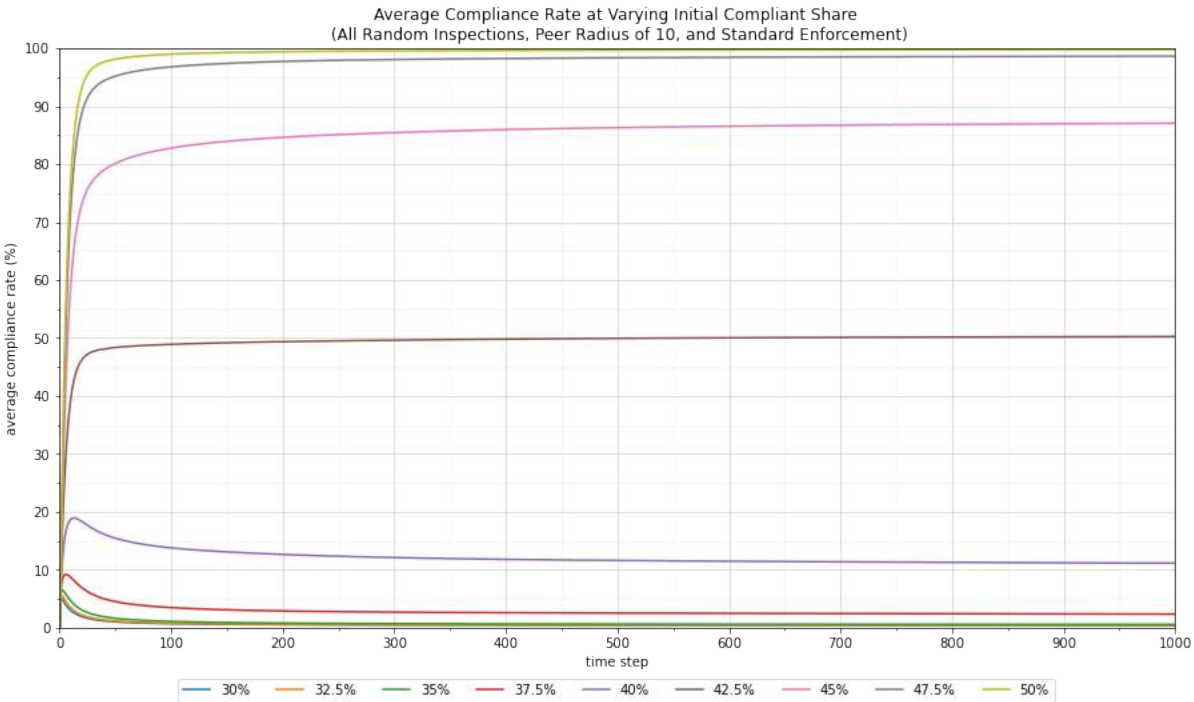

Figure A.1: Average compliance rate results with varying $\%_{init,comp}$ (SE).

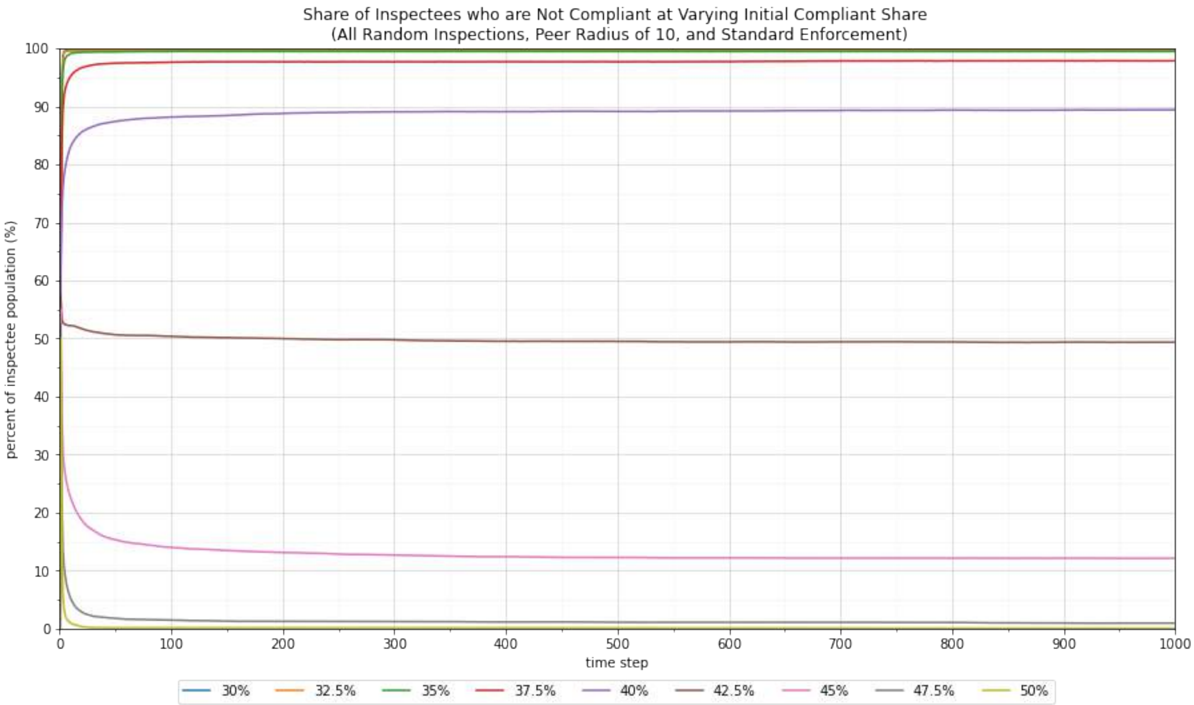

Figure A.2: Share of non-compliant inspectees with varying $\%_{init,comp}$ (SE).

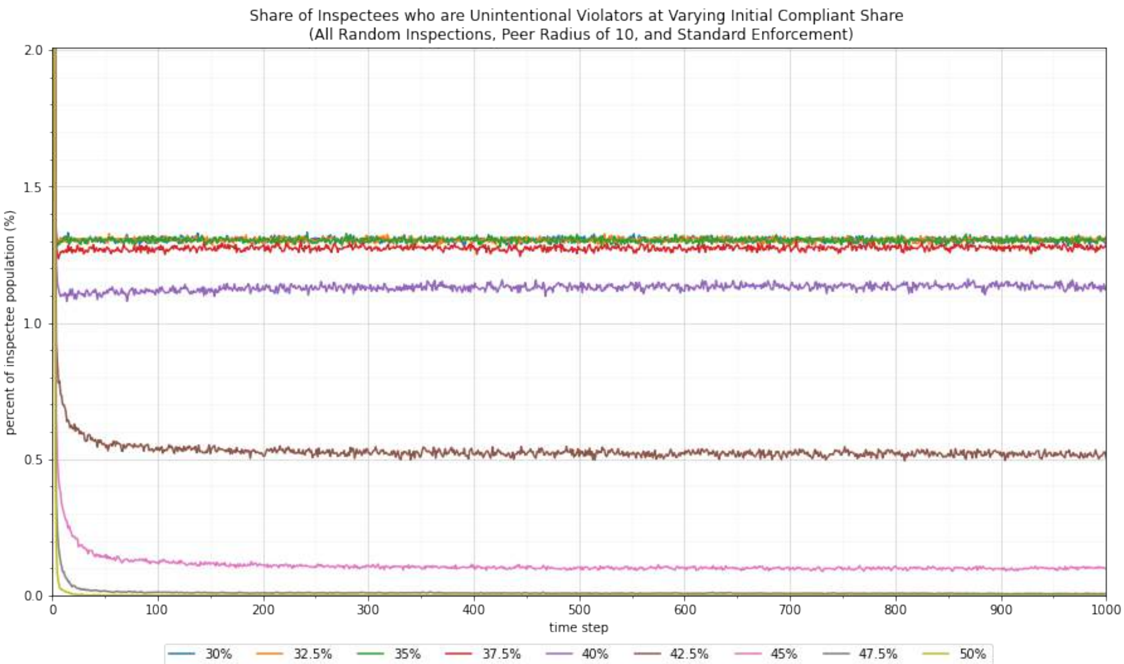

(a) Unintentional violators.

Figure A.3: Non-compliant inspectee breakdown with varying $\%_{init,comp}$ (SE).

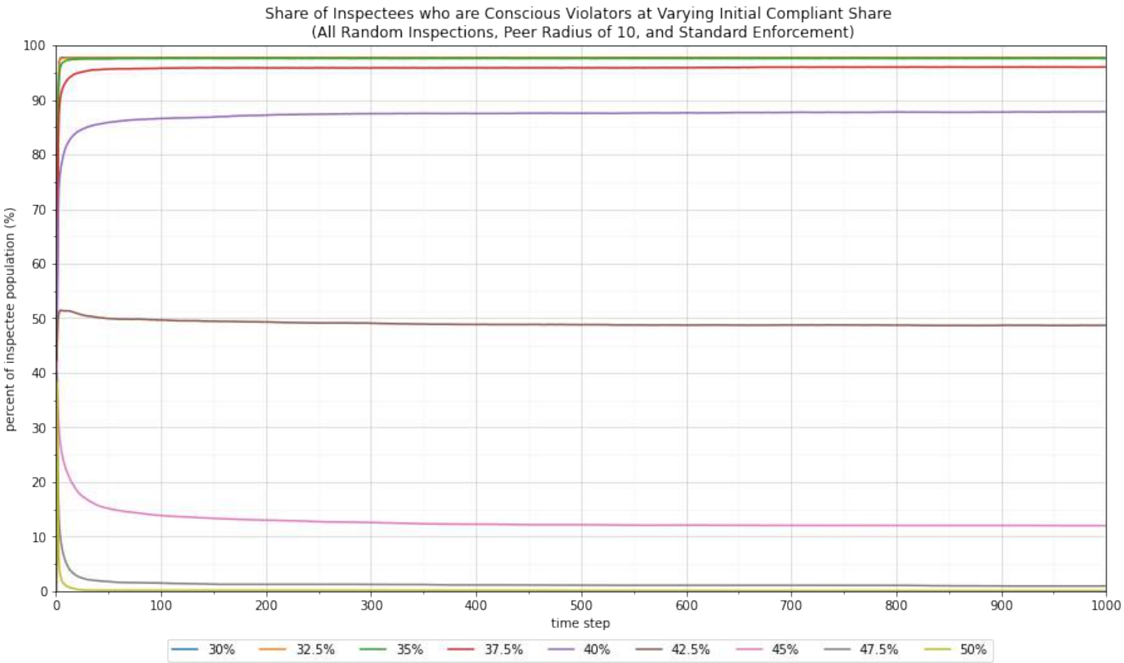

(b) Conscious violators.

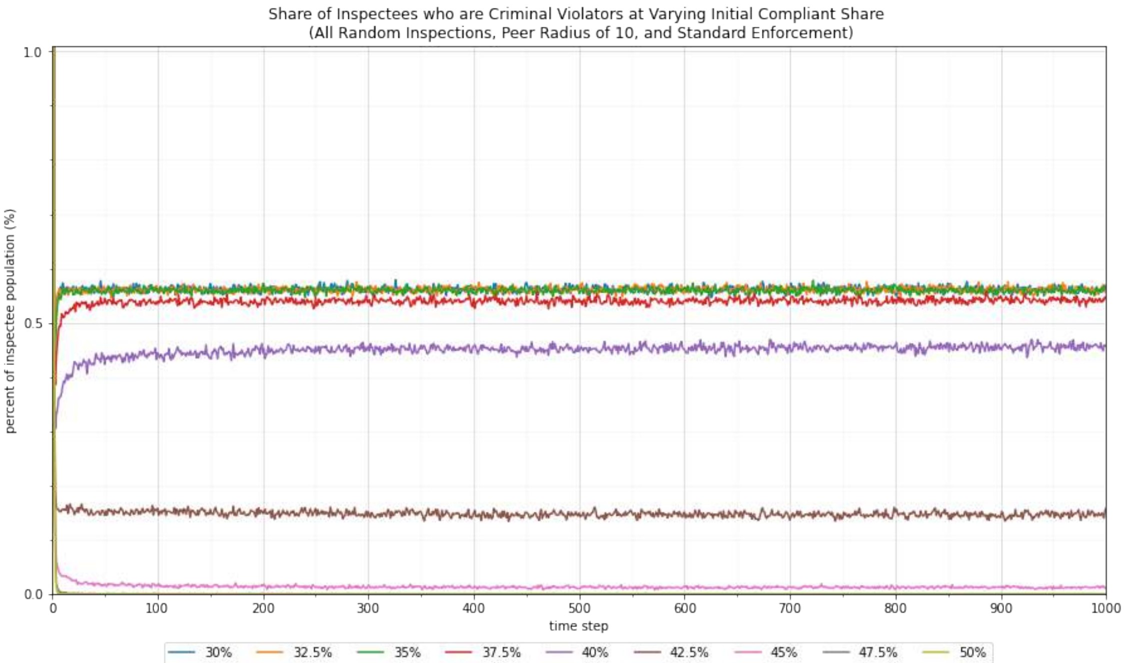

(c) Criminal violators.

Figure A.3: Non-compliant inspectee breakdown with varying $\%_{init,comp}$ (SE) (continued).

Responsive Enforcement

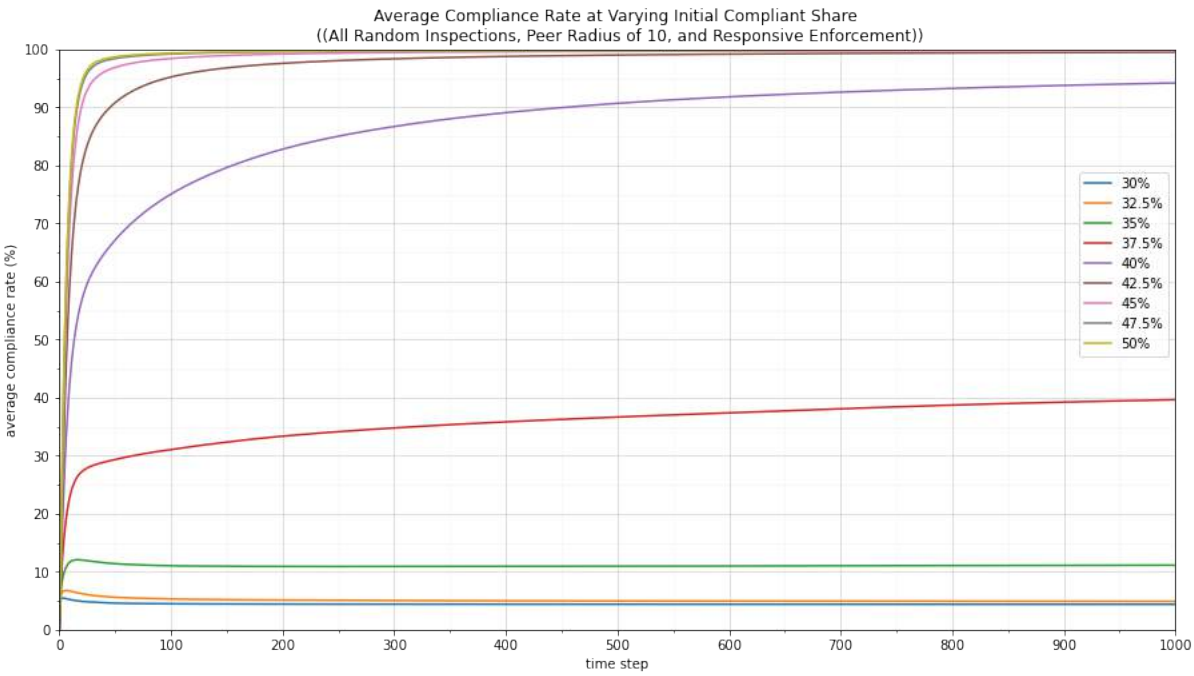

Figure A.4: Average compliance rate with varying $\%_{init,comp}$ (RE).

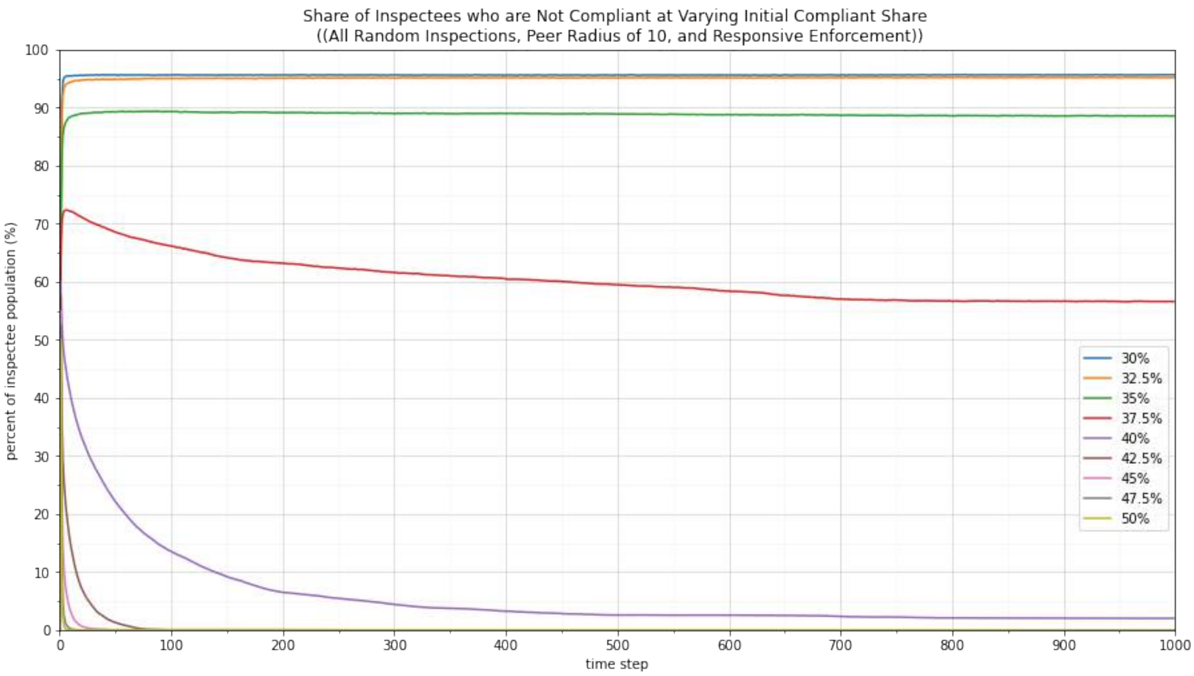

Figure A.5: Share of non-compliant inspectees with varying $\%_{init,comp}$ (RE).

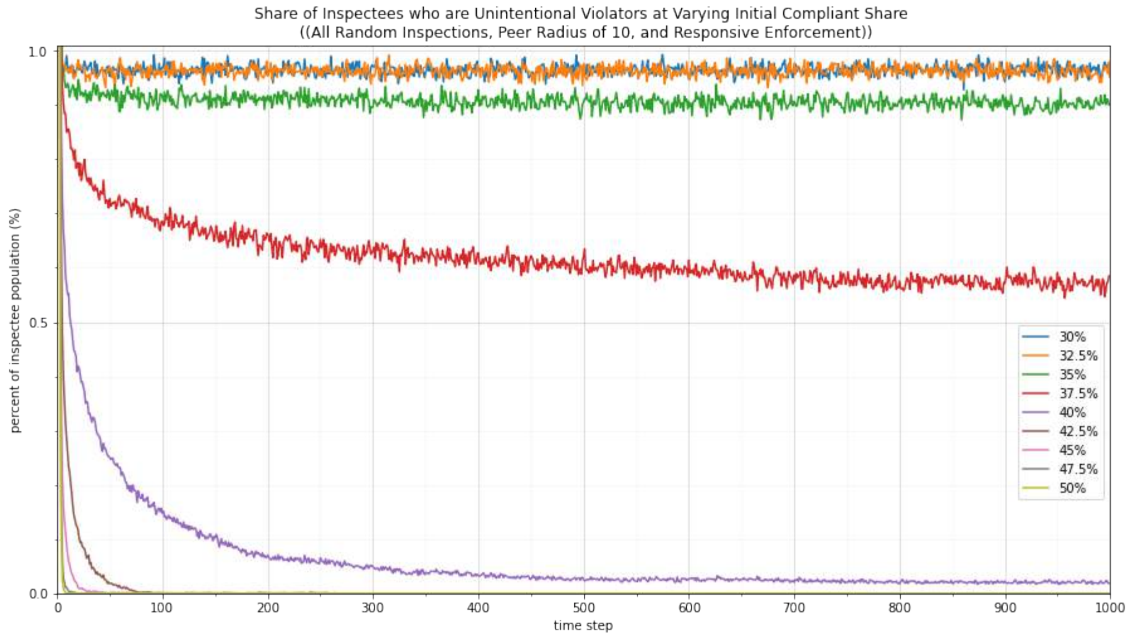

(a) Unintentional violators.

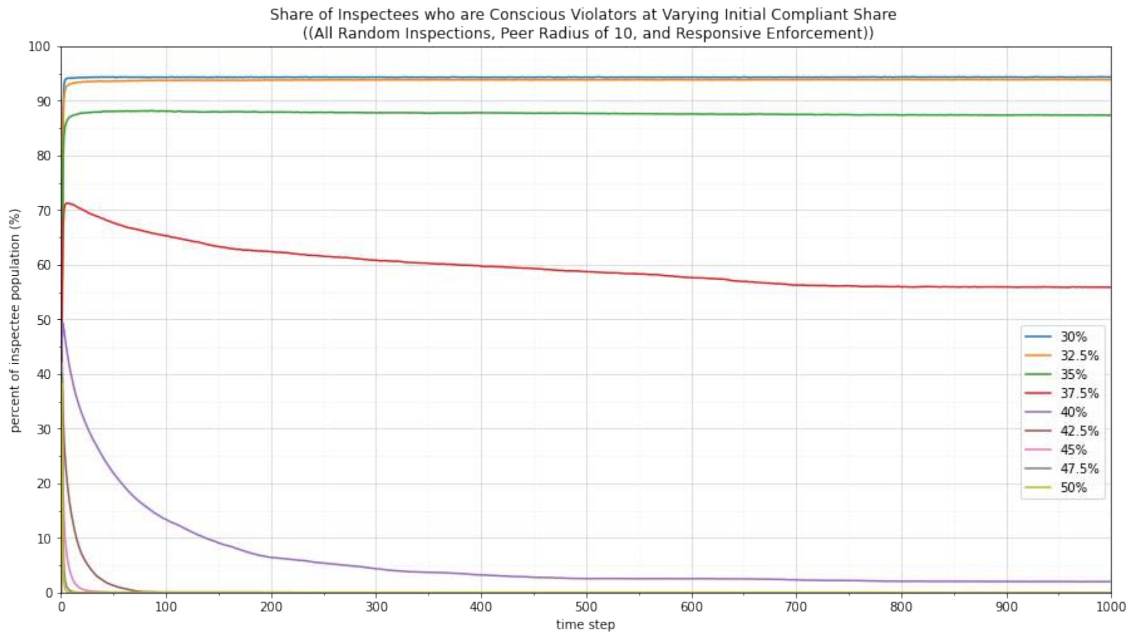

(b) Conscious violators.

Figure A.6: Non-compliant inspectee breakdown with varying $\%_{init,comp}$ (RE).

(c) Criminal violators.

Figure A.6: Non-compliant inspectee breakdown with varying $\%_{init,comp}$ (RE) (continued).

Path Dependency

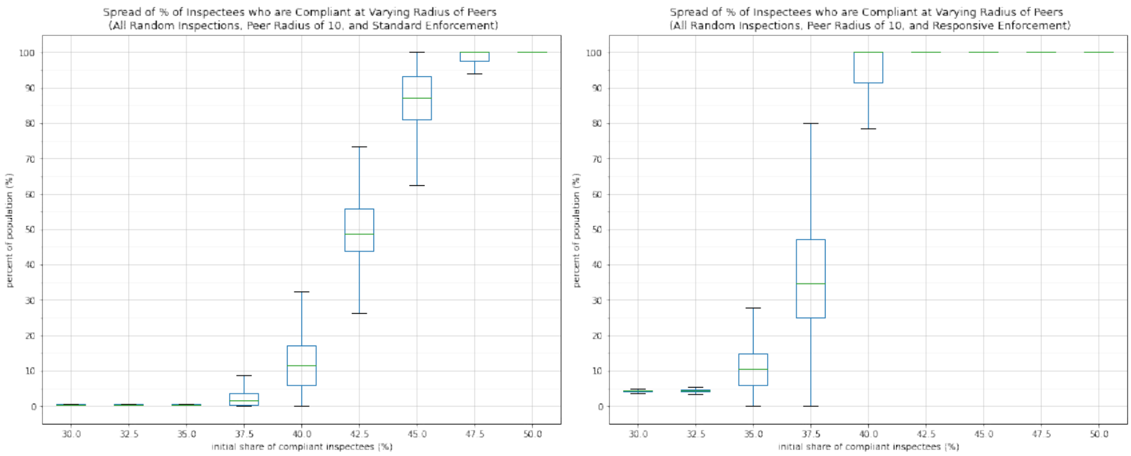

(a) Standard enforcement.

(b) Responsive enforcement.

Figure A.7: Spread of the share of compliant inspectees with varying $\%_{init,comp}$.

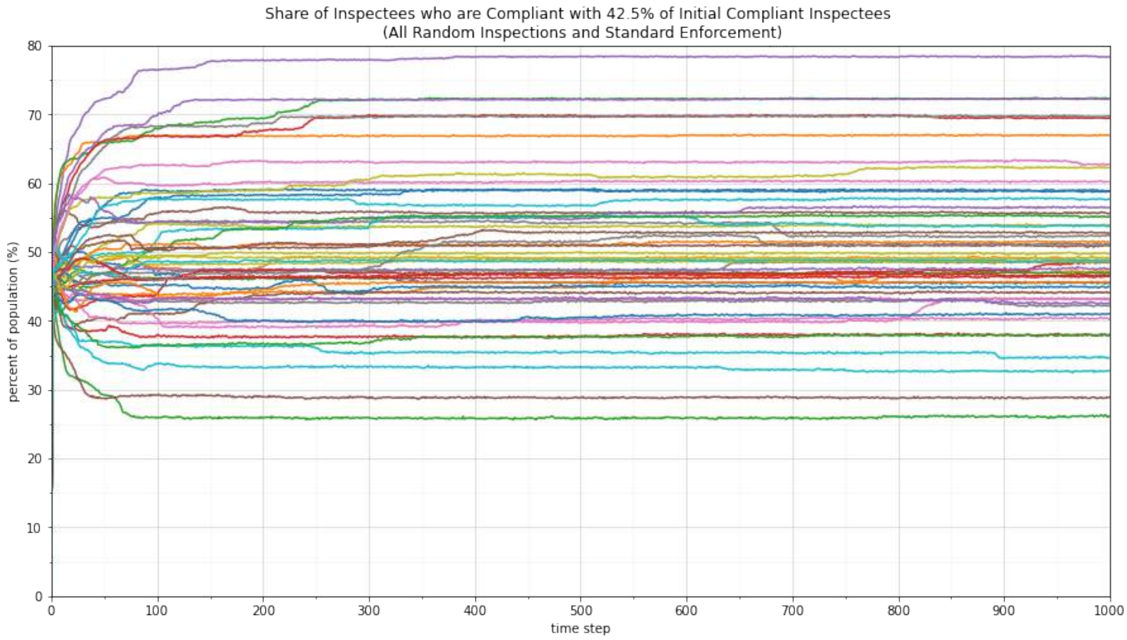

(a) Compliant inspectees.

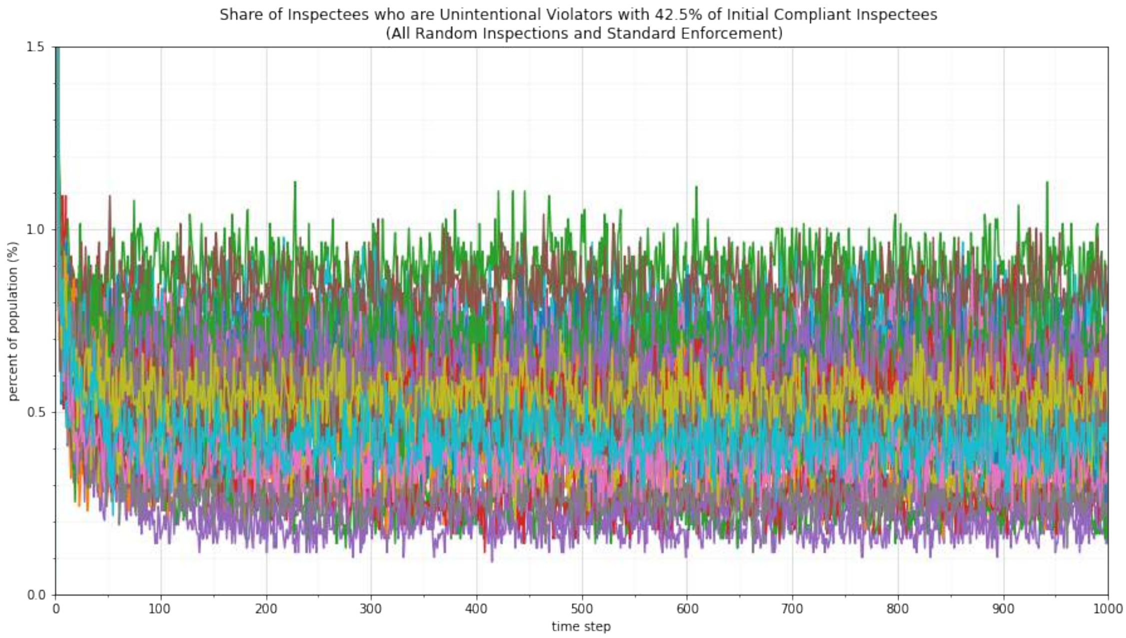

(b) Unintentional violators.

Figure A.8: Compliance level breakdown with $\%_{init,comp}$=42.5% (SE).

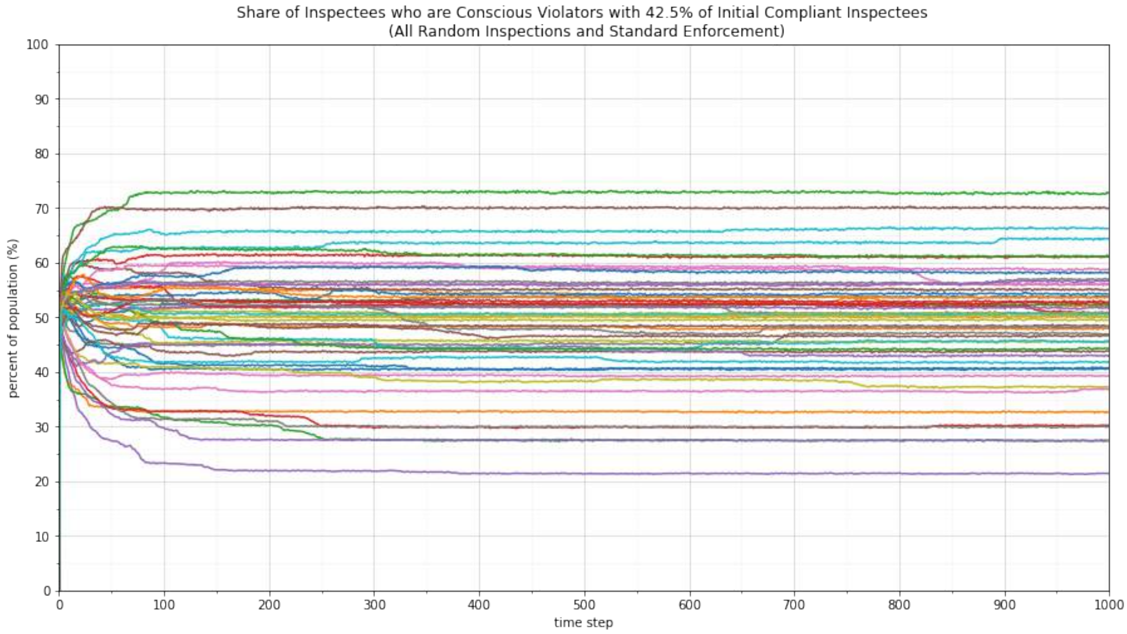

(c) Conscious violators.

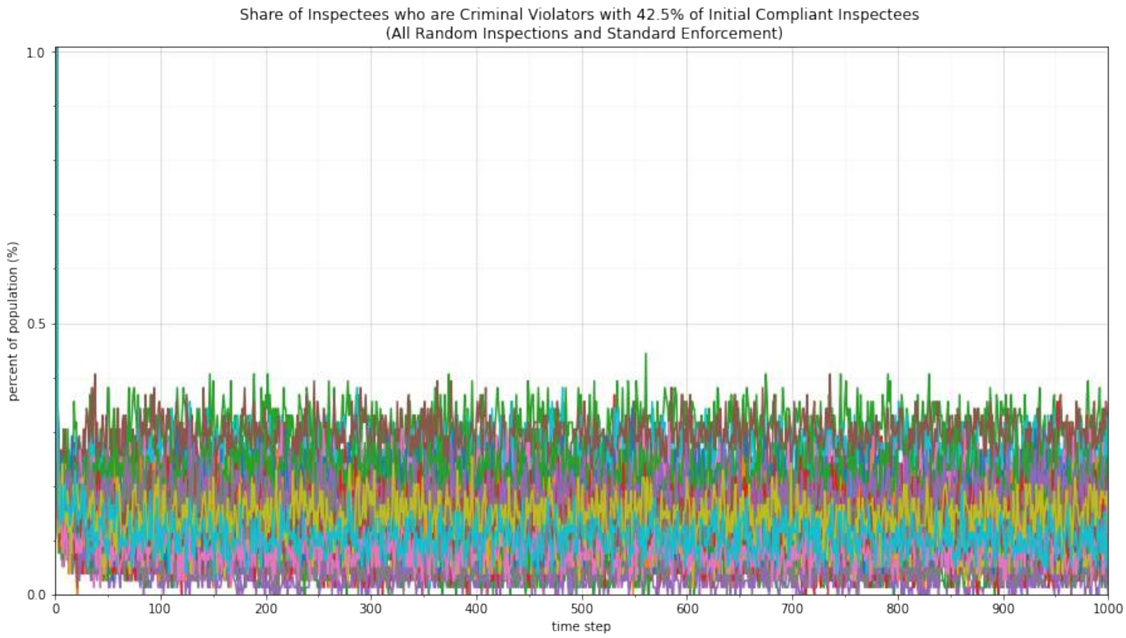

(d) Criminal violators.

Figure A.8: Compliance level breakdown with $\%_{init,comp}$=42.5% (SE) (continued).

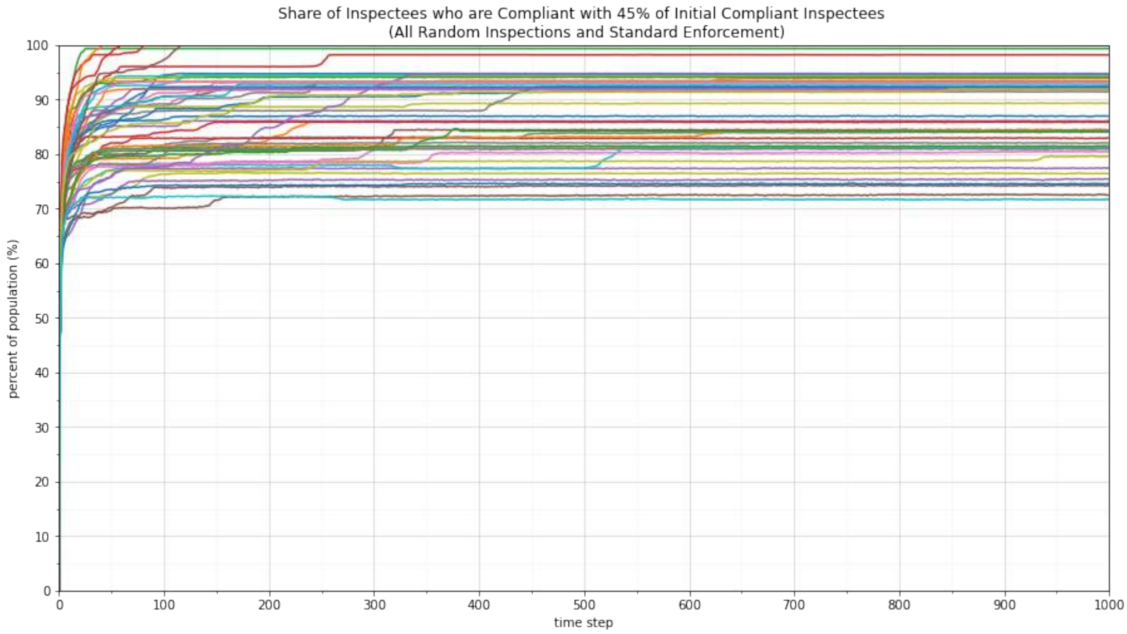

(a) Compliant inspectees.

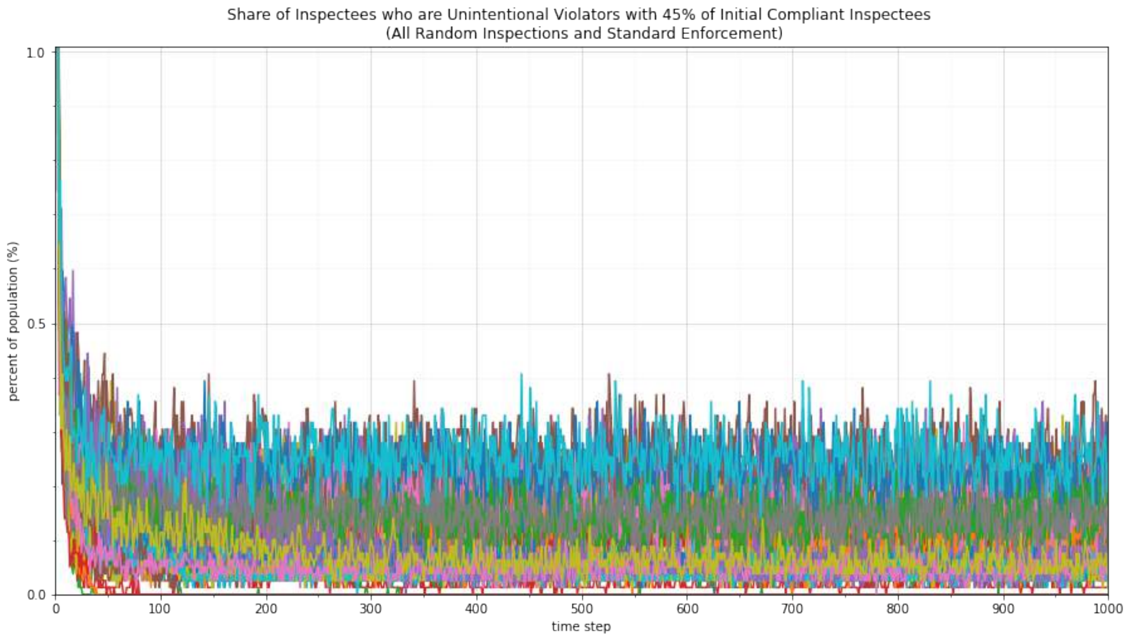

(b) Unintentional violators.

Figure A.9: Compliance level breakdown with $\%_{init,comp}$=45% (SE).

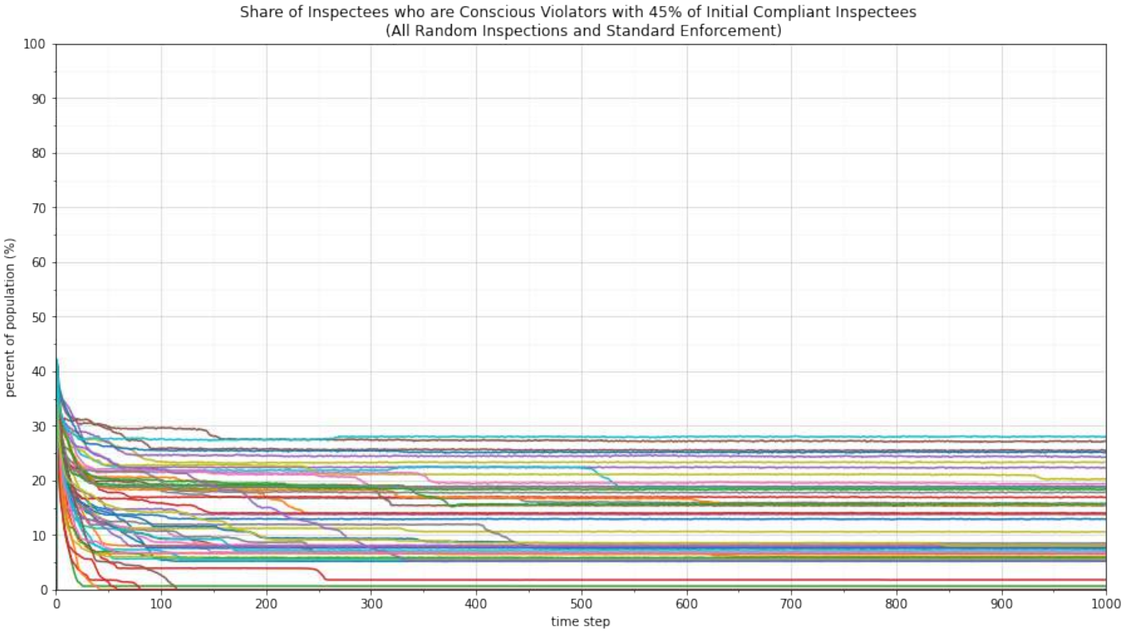

(c) Conscious violators.

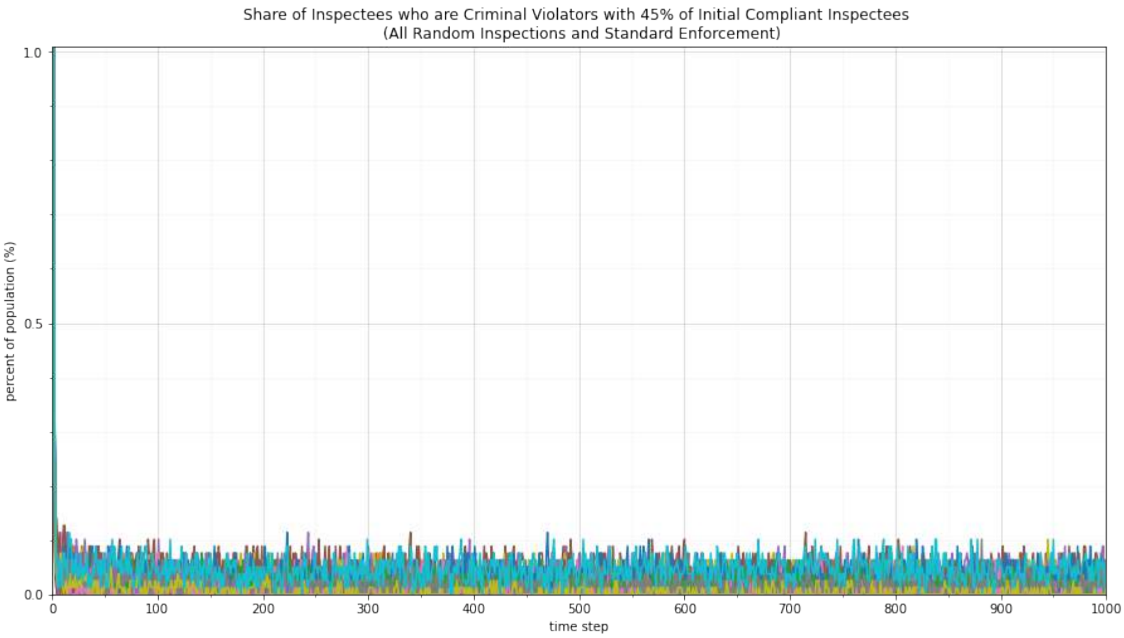

(d) Criminal violators.

Figure A.9: Compliance level breakdown with $\%_{init,comp}$=45% (SE) (continued).

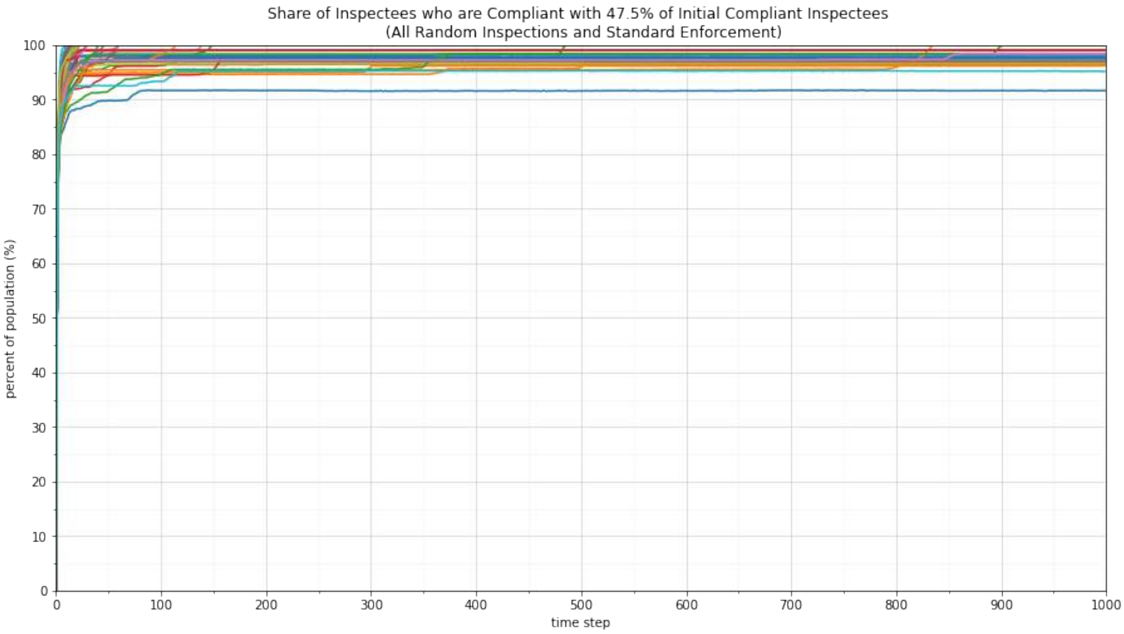

(a) Compliant inspectees.

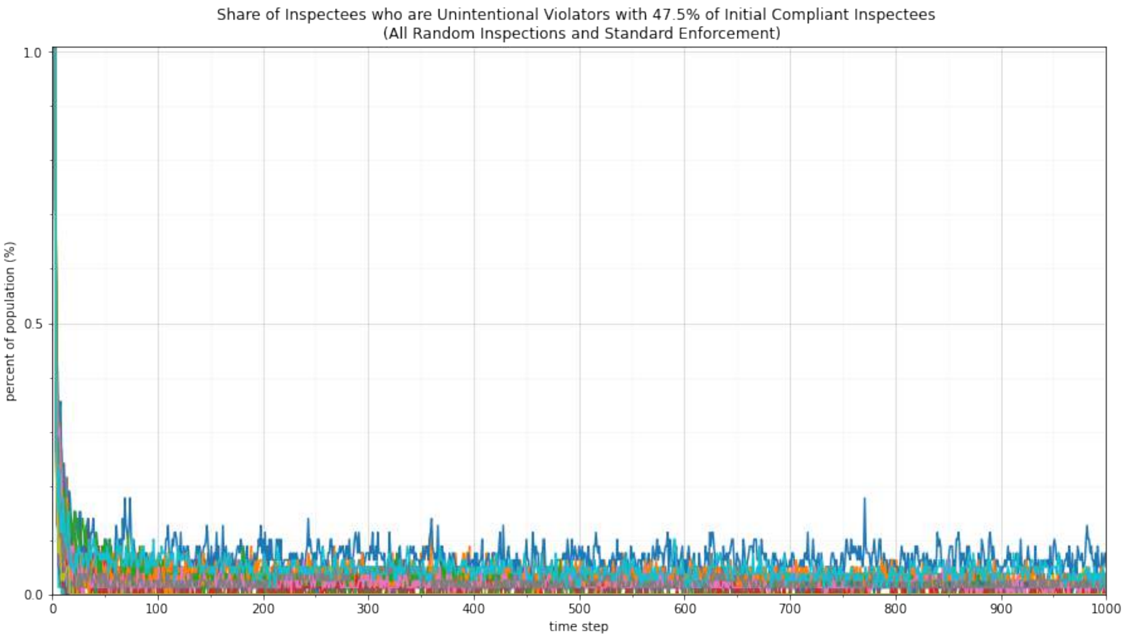

(b) Unintentional violators.

Figure A.10: Compliance level breakdown with $\%_{init,comp}$=47.5% (SE).

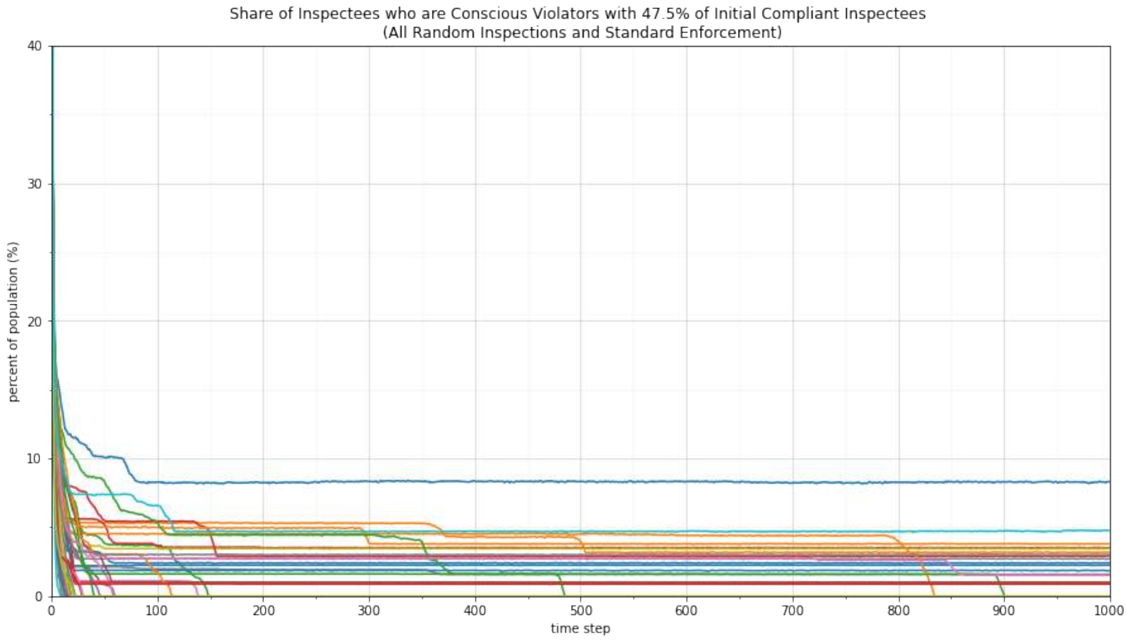

(c) Conscious violators.

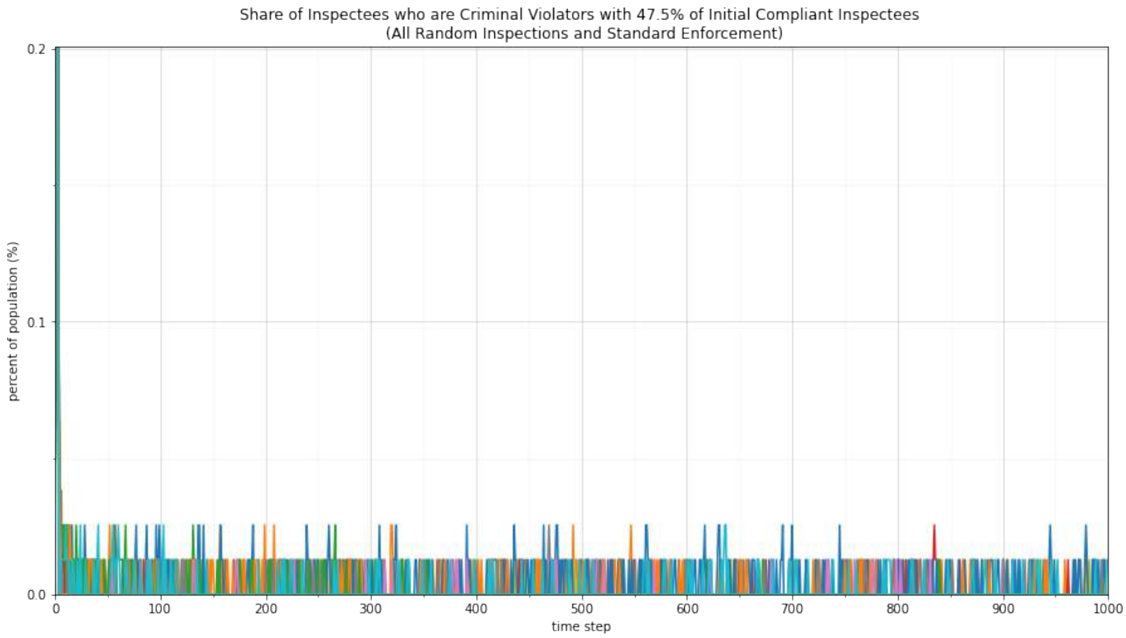

(d) Criminal violators.

Figure A.10: Compliance level breakdown with $\%_{init,comp}$=47.5% (SE) (continued).

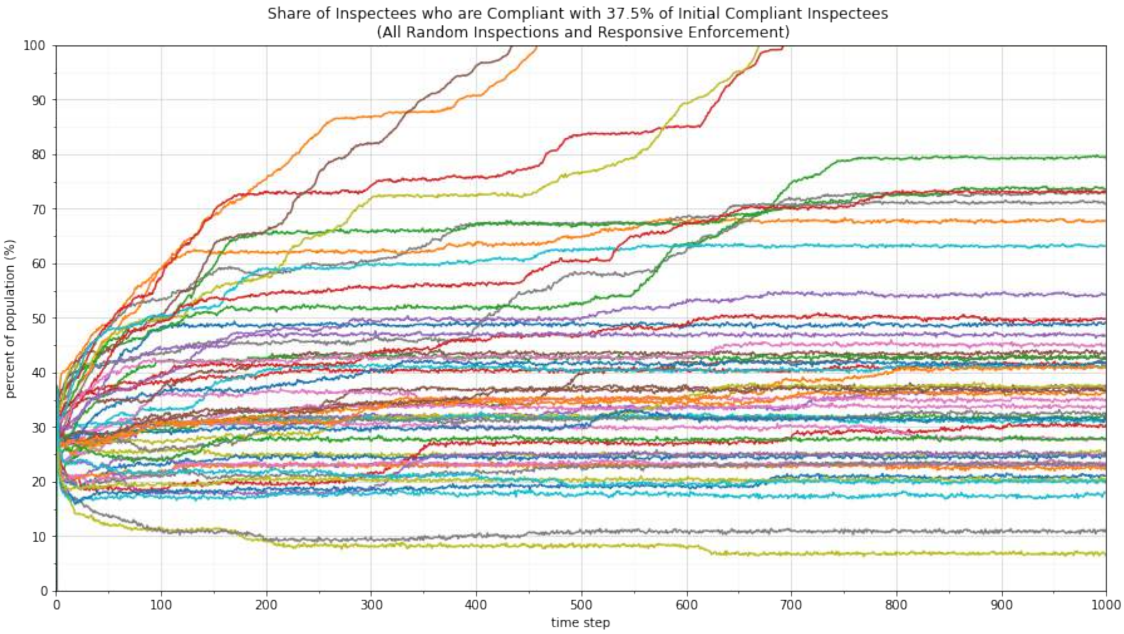

(a) Compliant inspectees.

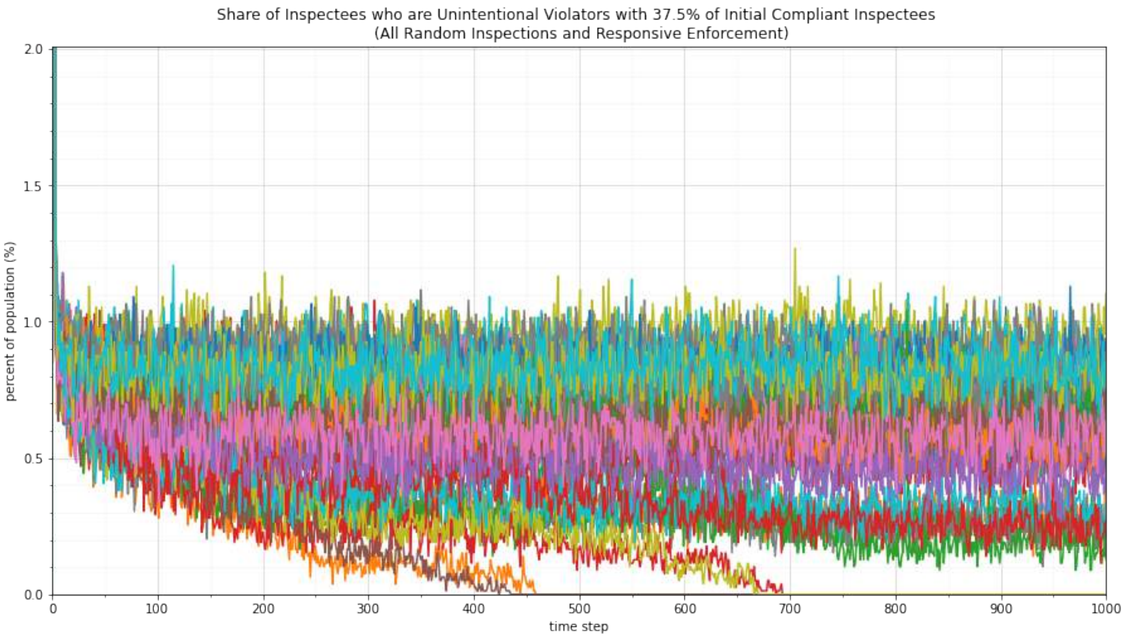

(b) Unintentional violators.

Figure A.11: Compliance level breakdown with $\%_{init,comp}$=37.5% (RE).

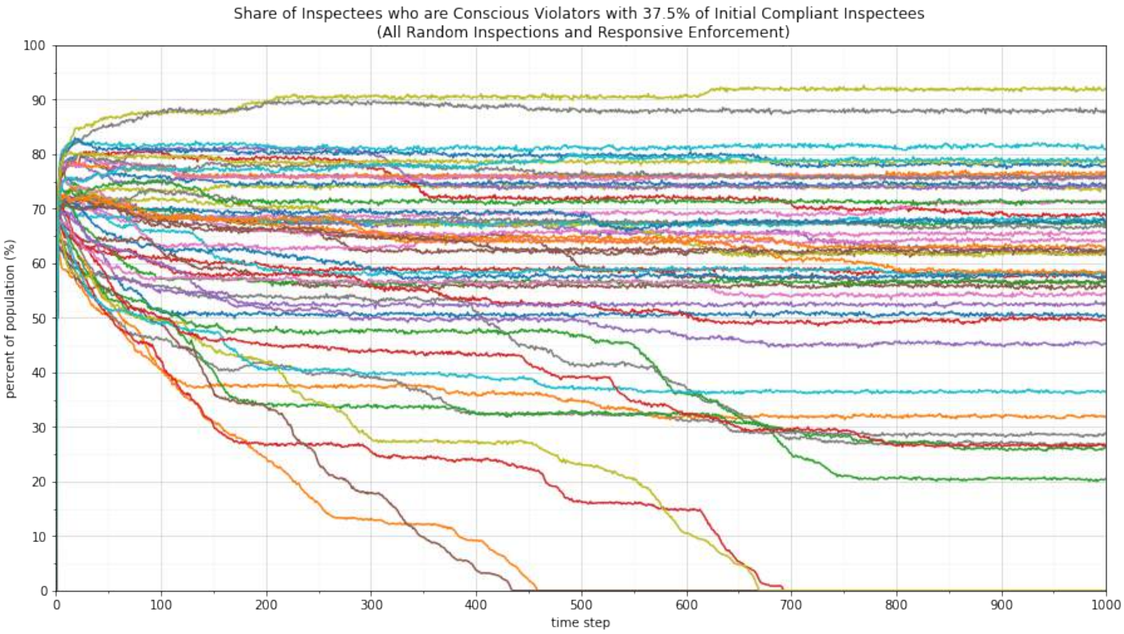

(c) Conscious violators.

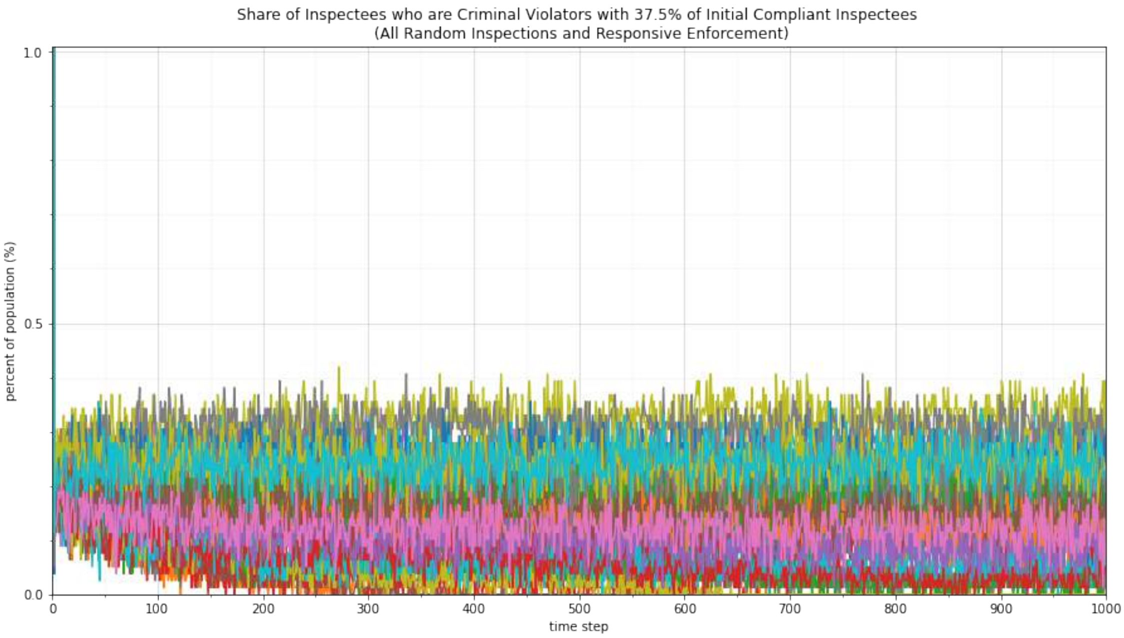

(d) Criminal violators.

Figure A.11: Compliance level breakdown with $\%_{init,comp}$=37.5% (RE) (continued).

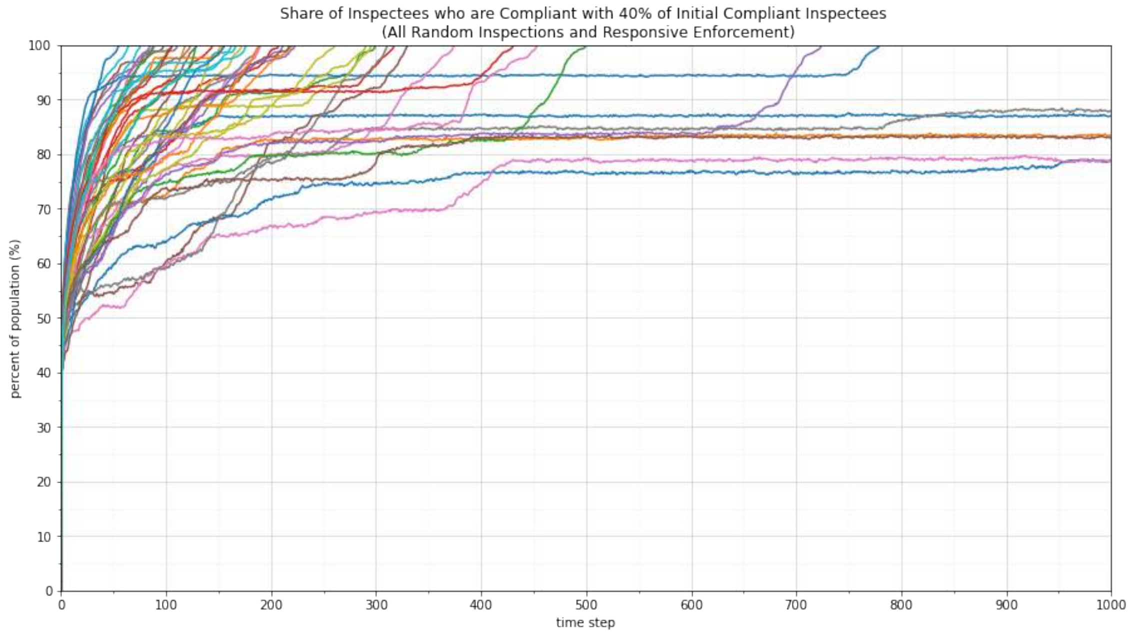

(a) Compliant inspectees.

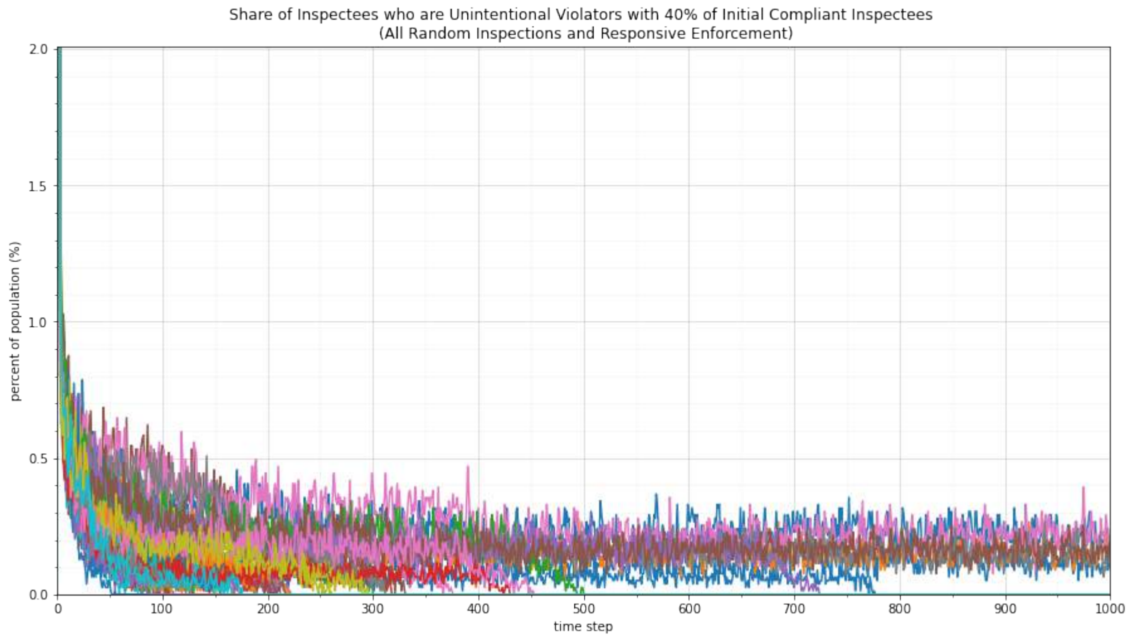

(b) Unintentional violators.

Figure A.12: Compliance level breakdown with $\%_{init,comp}$=40% (RE).

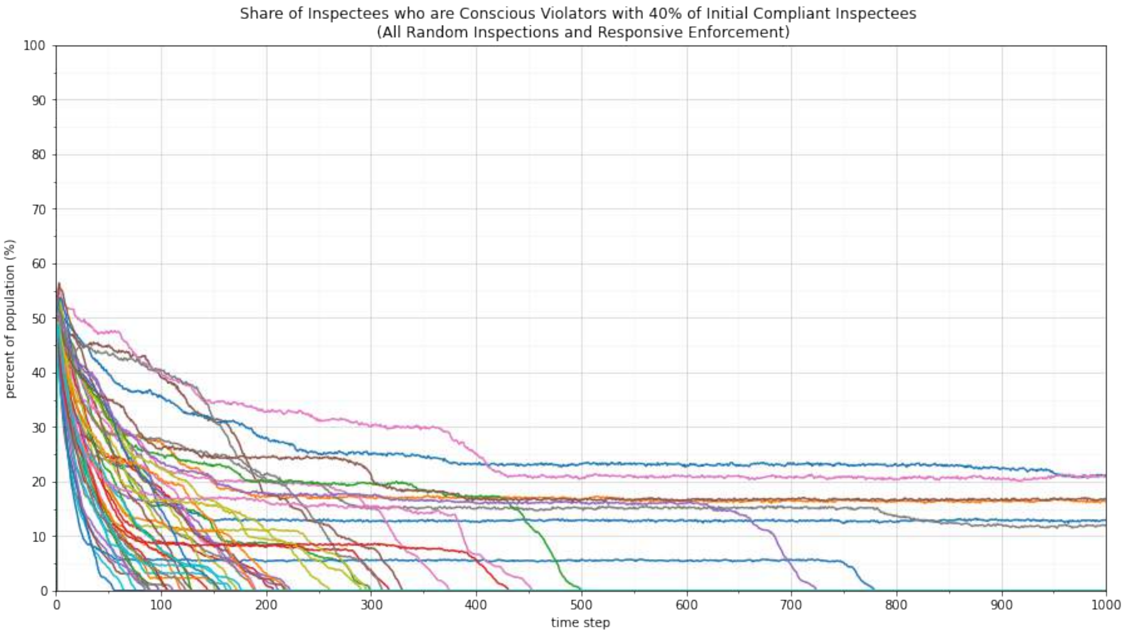

(c) Conscious violators.

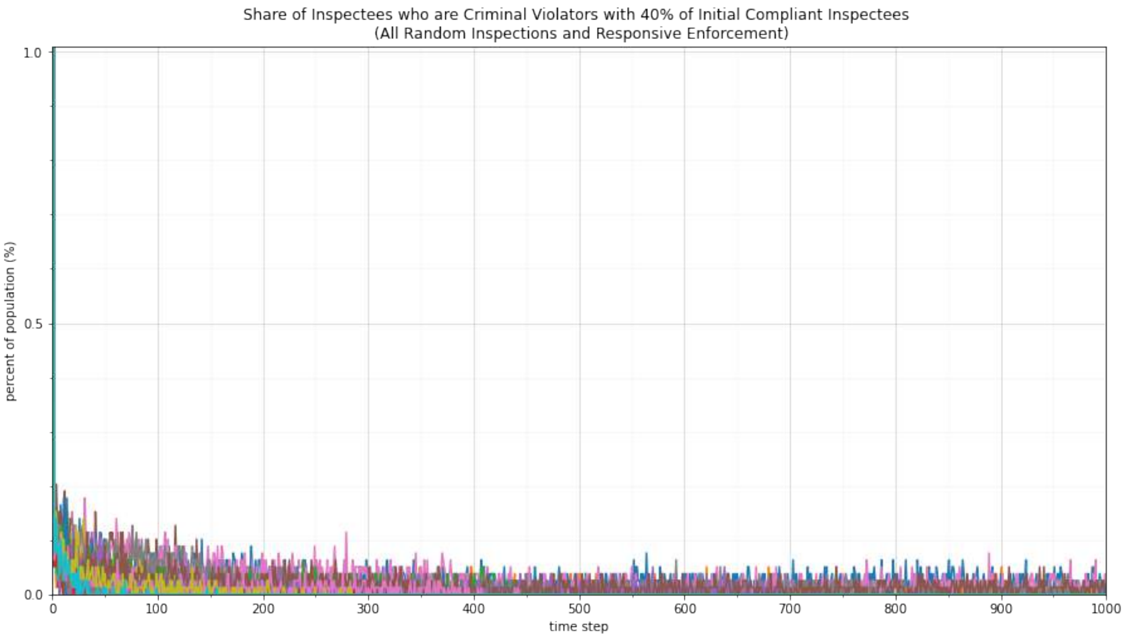

(d) Criminal violators.

Figure A.12: Compliance level breakdown with $\%_{init,comp}$=40% (RE) (continued).

A.1.2 Varying Percentage of Random Inspections ($\%_{rand}$)

A.1.2.1 Mixed Random and Risk-based Inspection Strategy based on Non-compliant Record (MRRbNC)

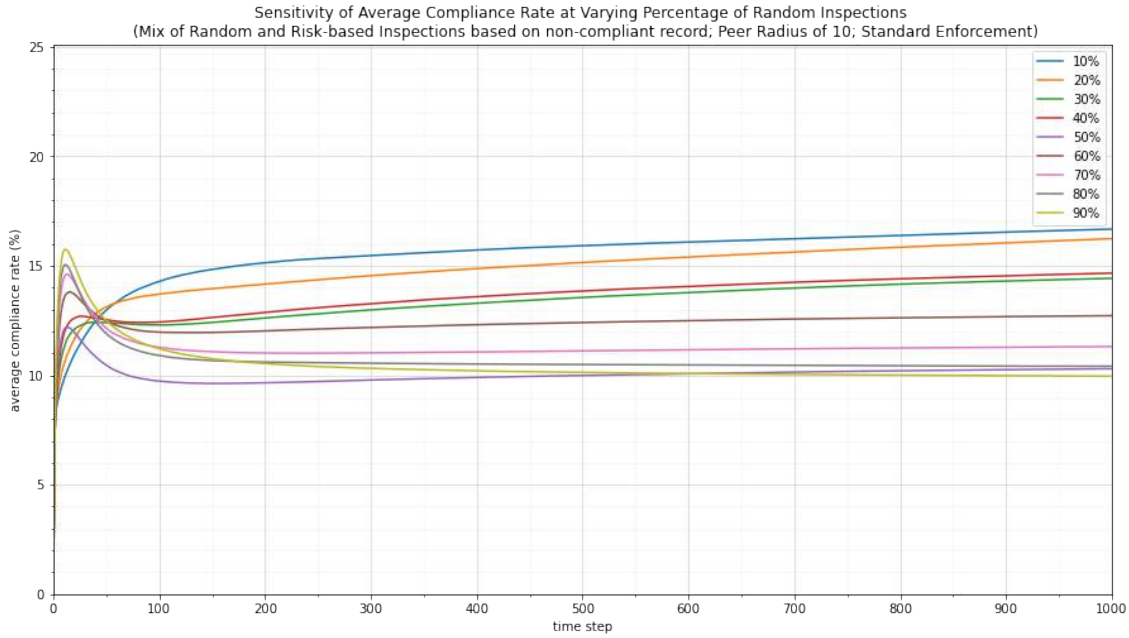

(a) Standard enforcement.

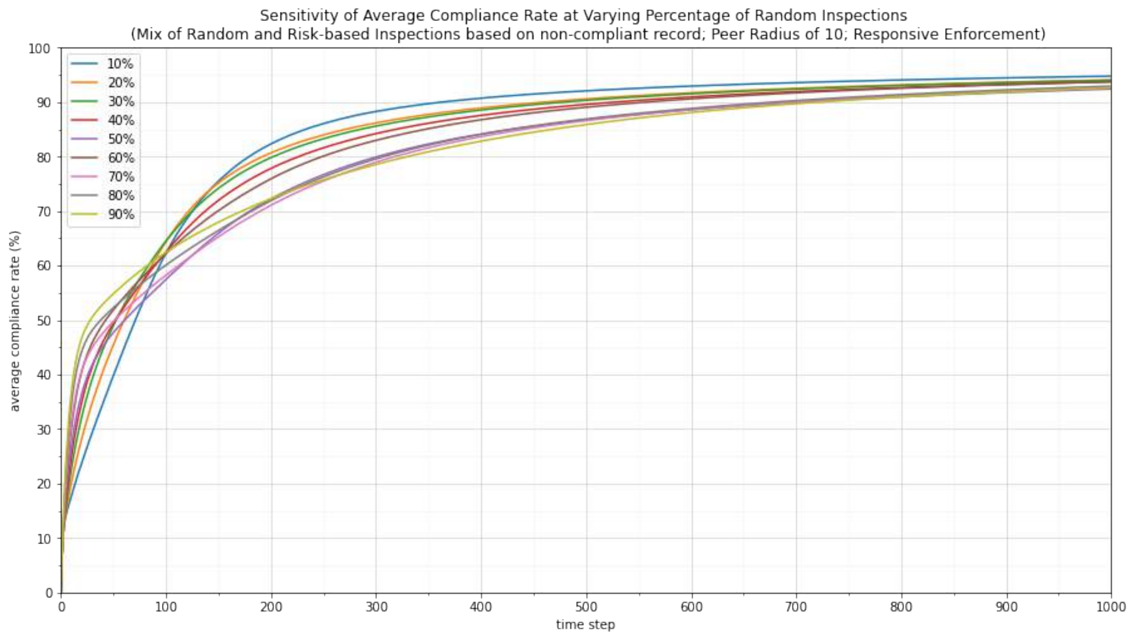

(b) Responsive enforcement.

Figure A.13: Sensitivity of the compliance level breakdown with varying $\%_{rand}$ (RE, MRRbNC).

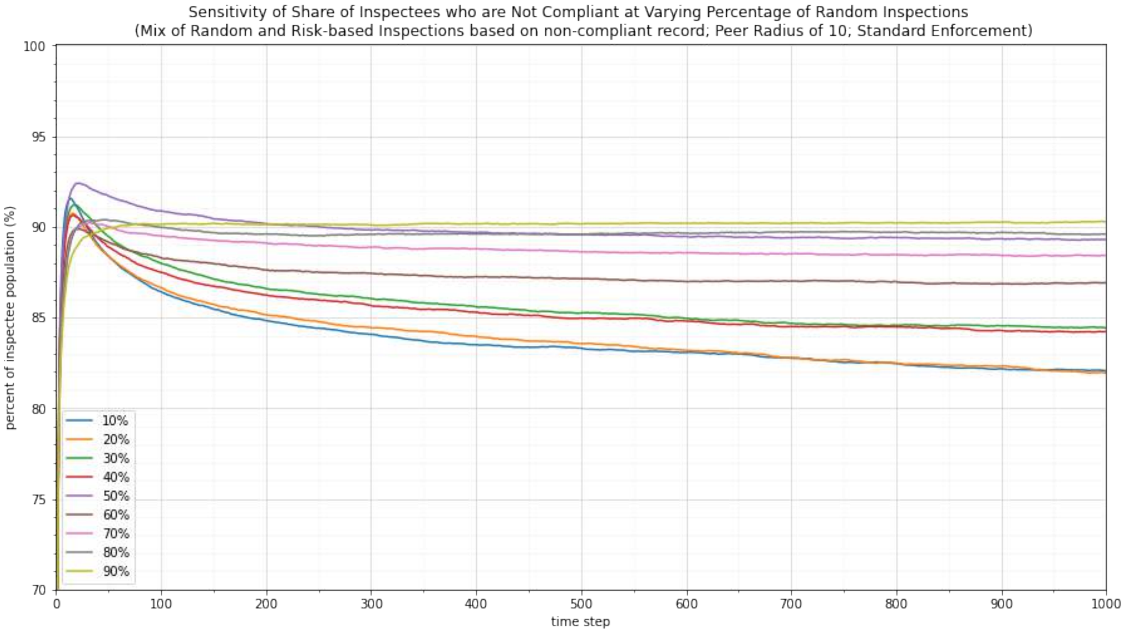

(a) Standard enforcement.
For a detailed breakdown by compliance level, see Figure A.15.

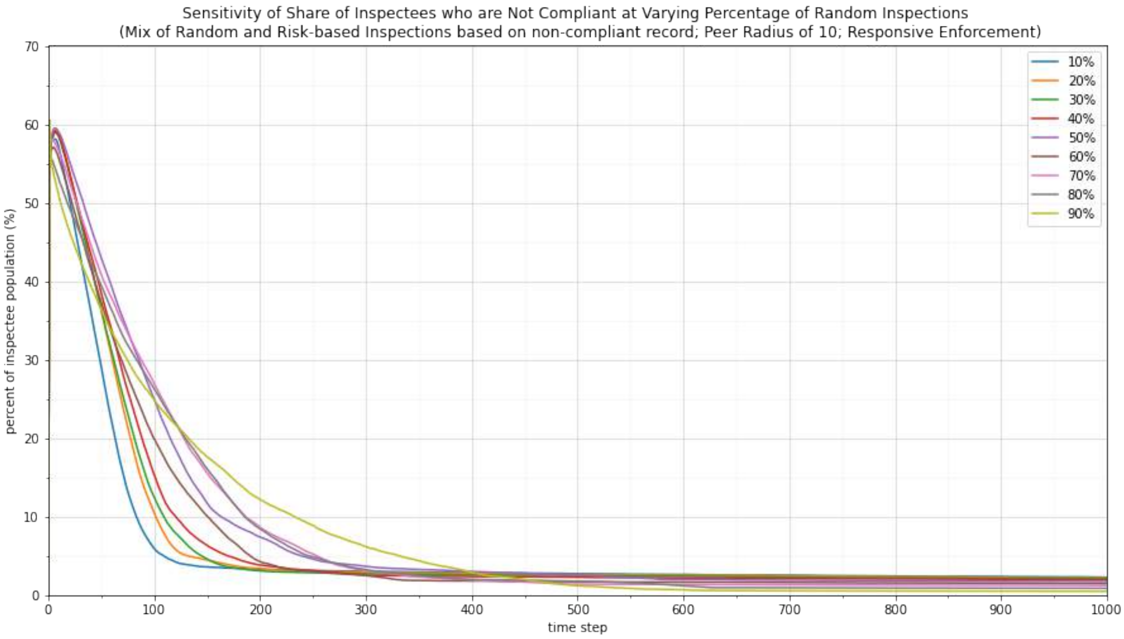

(b) Responsive enforcement.
For a detailed breakdown by compliance level, see Figure A.16.

Figure A.14: Share of non-compliant inspectees with varying$\%_{rand}$ (MRRbNC).

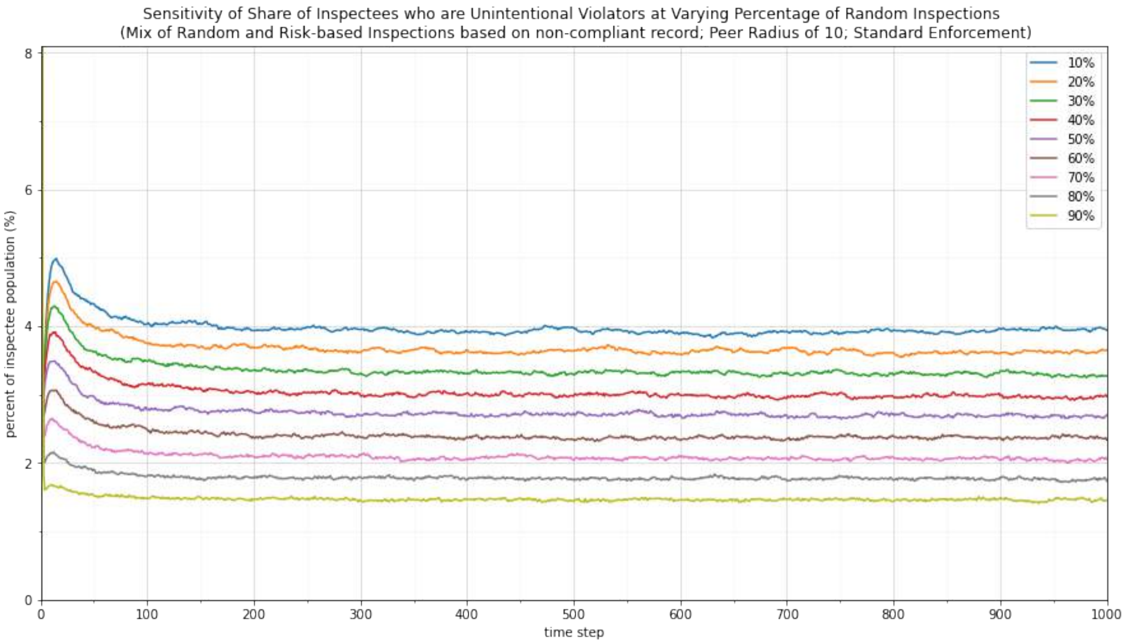

(a) Unintentional violators.

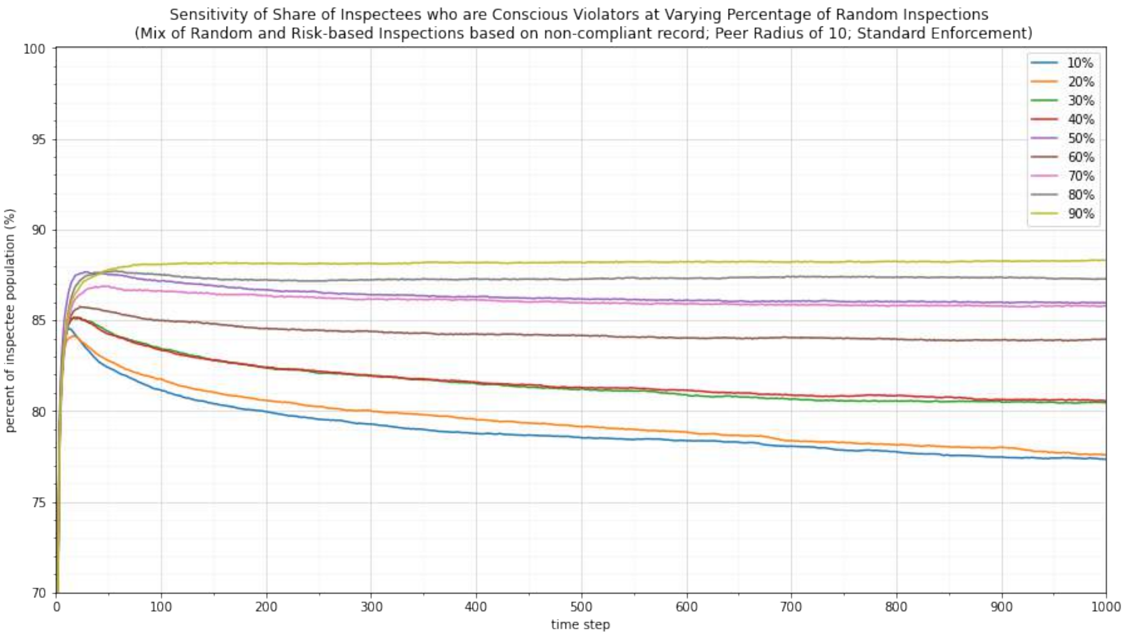

(b) Conscious violators.

Figure A.15: Non-compliant inspectee breakdown with varying $\%_{rand}$ (SE)

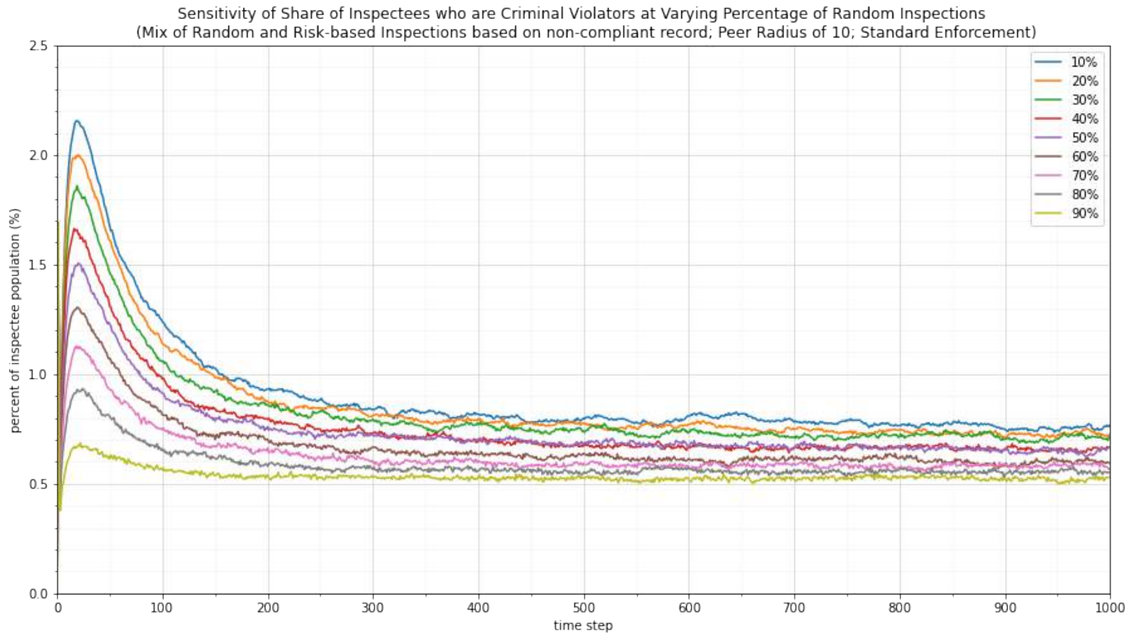

(c) Criminal violators.

Figure A.15: Non-compliant inspectee breakdown with varying $\%_{rand}$ (SE) (continued)..

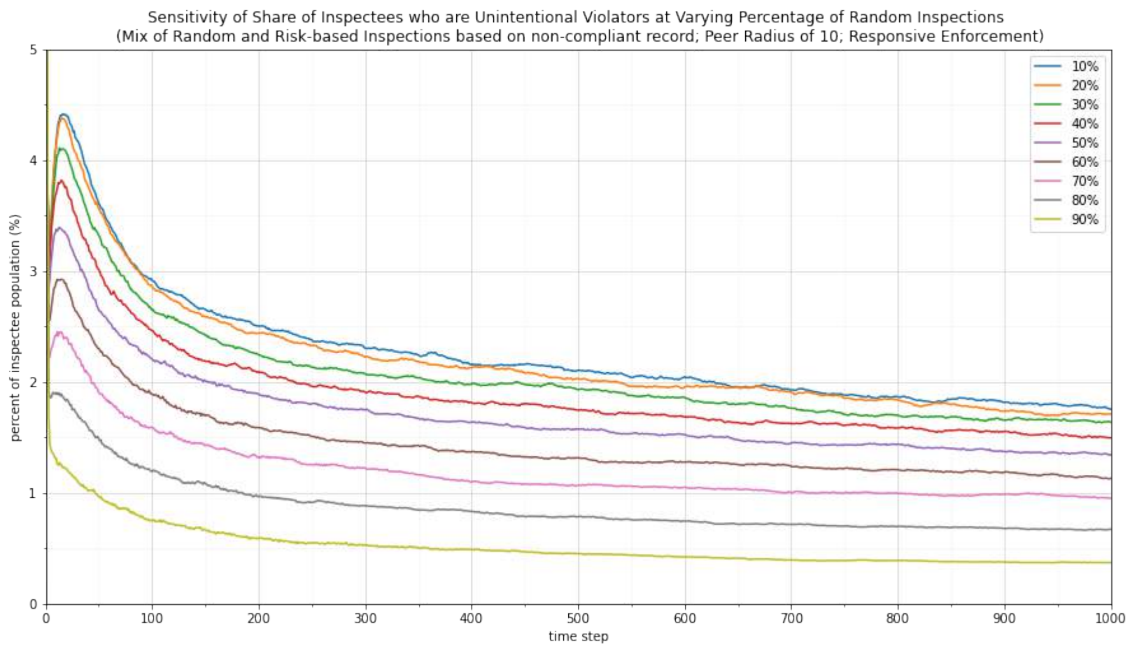

(a) Unintentional violators.

Figure A.16: Non-compliant inspectee breakdown with varying $\%_{rand}$ (RE).

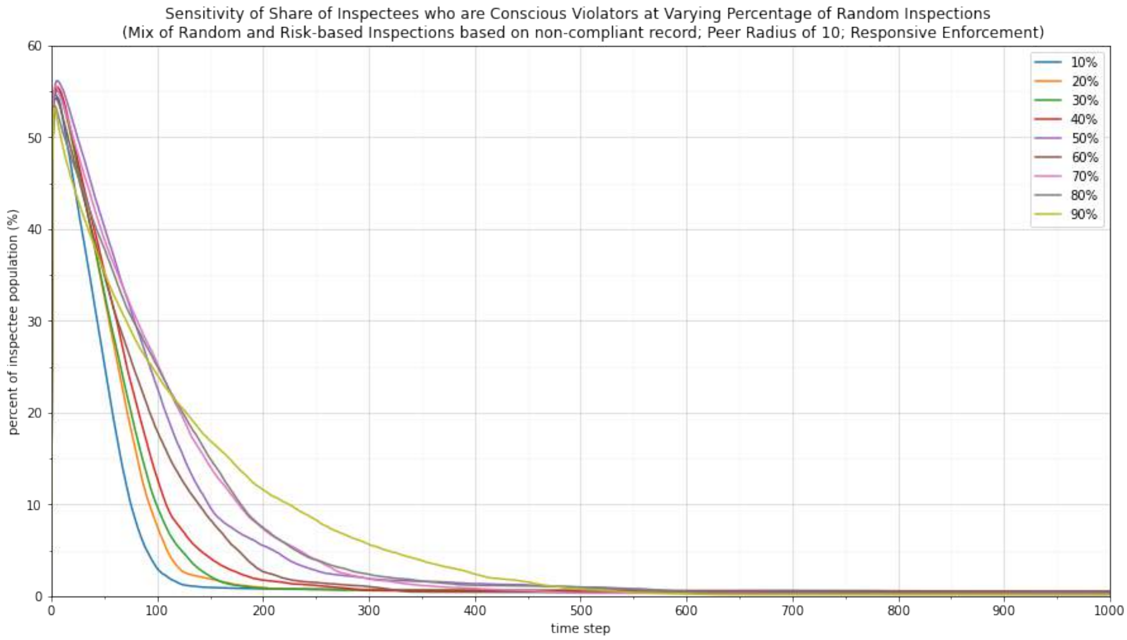

(b) Conscious violators.

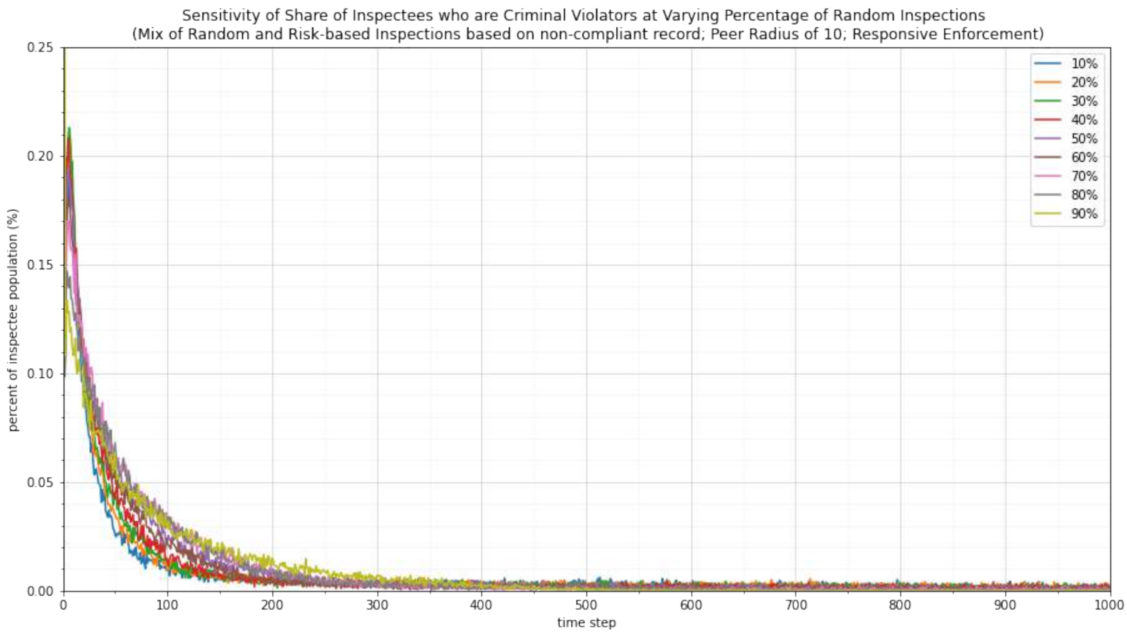

(c) Criminal violators.

Figure A.16: Non-compliant inspectee breakdown with varying $\%_{rand}$ (RE) (continued).

Path Dependency

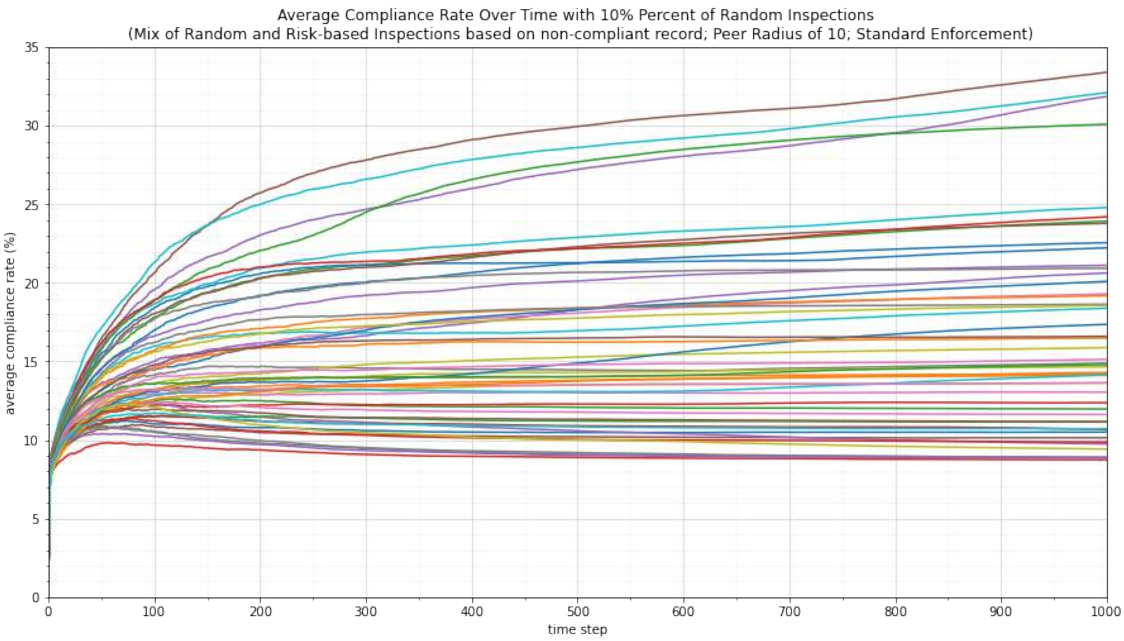

(a) 10% random inspections (SE).

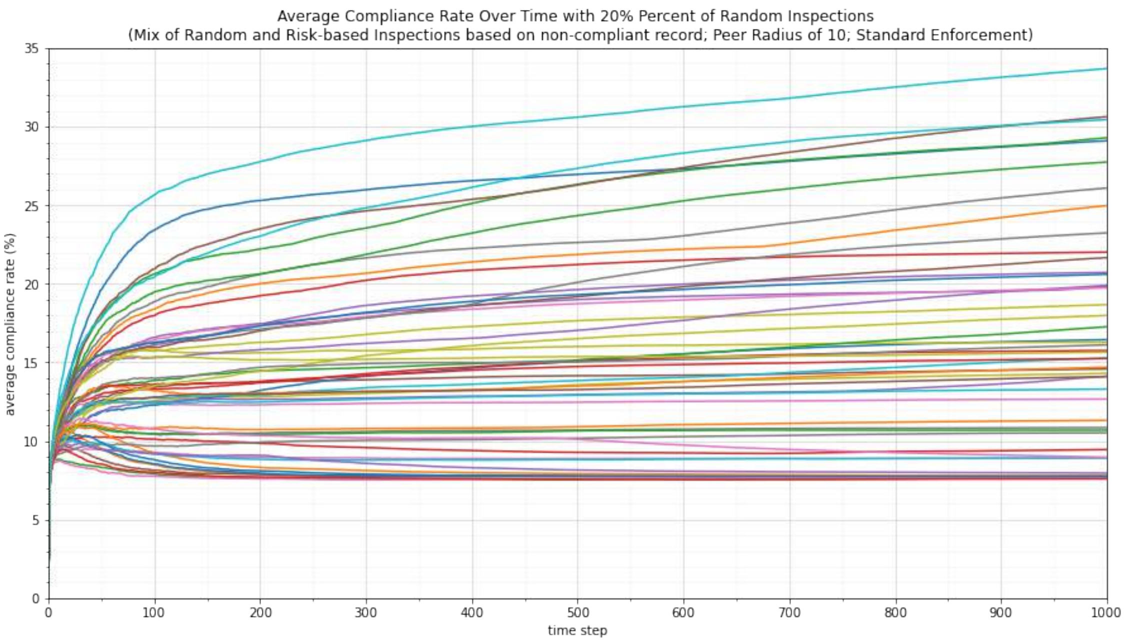

(b) 20% random inspections (SE).

Figure A.17: Average compliance rate results with varying $\%_{rand}$ (SE, MRRbNC).

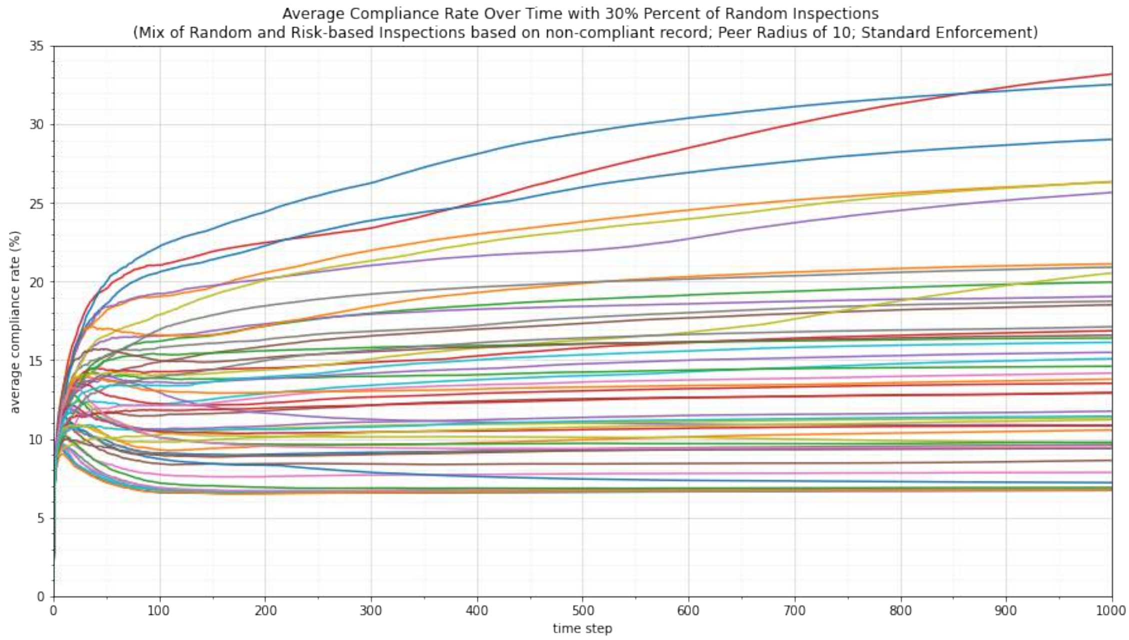

(c) 30% random inspections (SE).

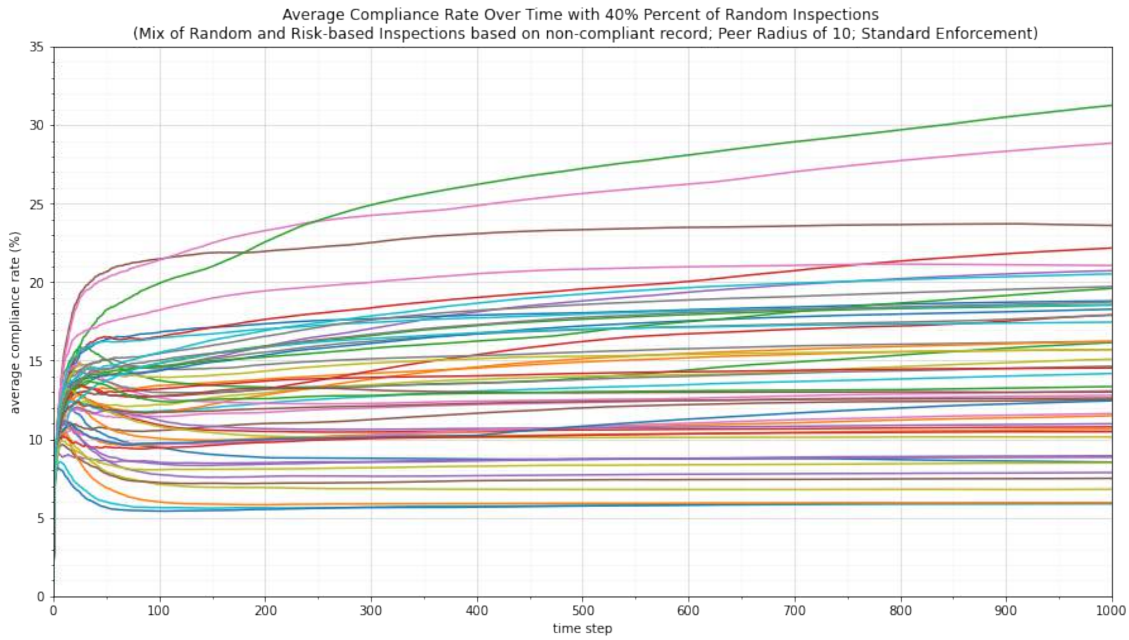

(d) 40% random inspections (SE).

Figure A.17: Average compliance rate results with varying $\%_{rand}$ (SE, MRRbNC, continued).

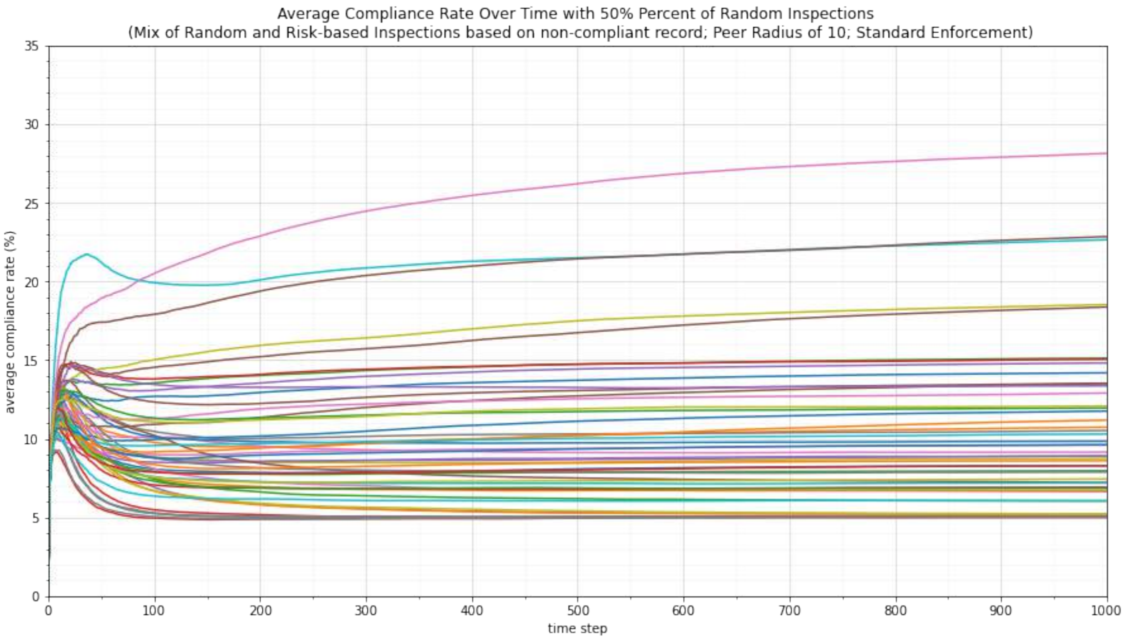

(e) 50% random inspections (SE).

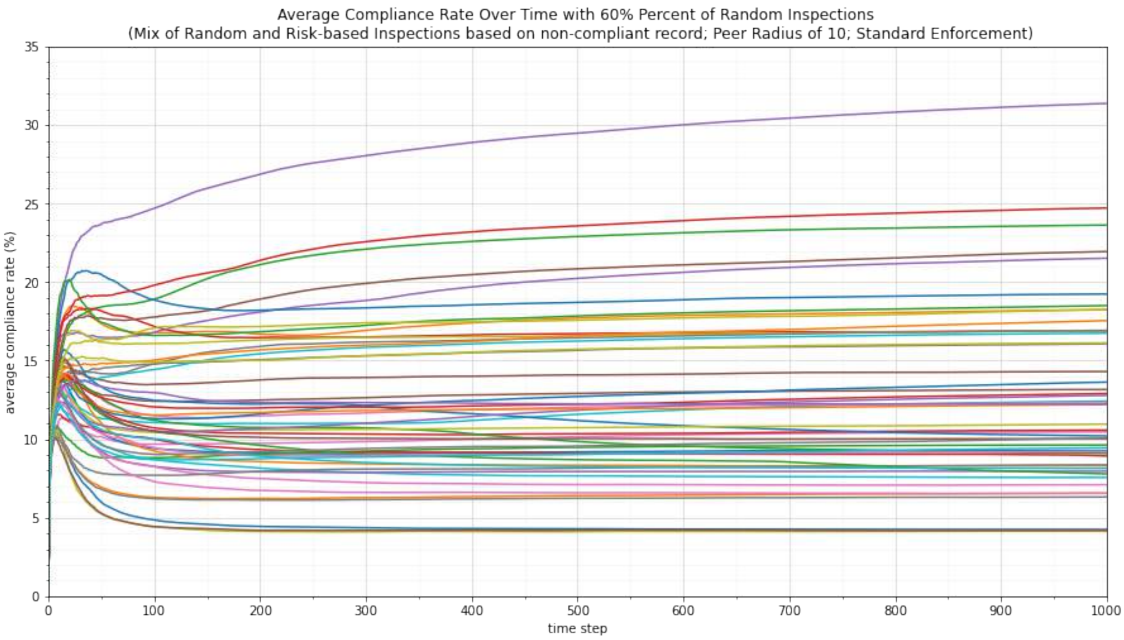

(f) 60% random inspections (SE).

Figure A.17: Average compliance rate results with varying $\%_{rand}$ (SE, MRRbNC, continued).

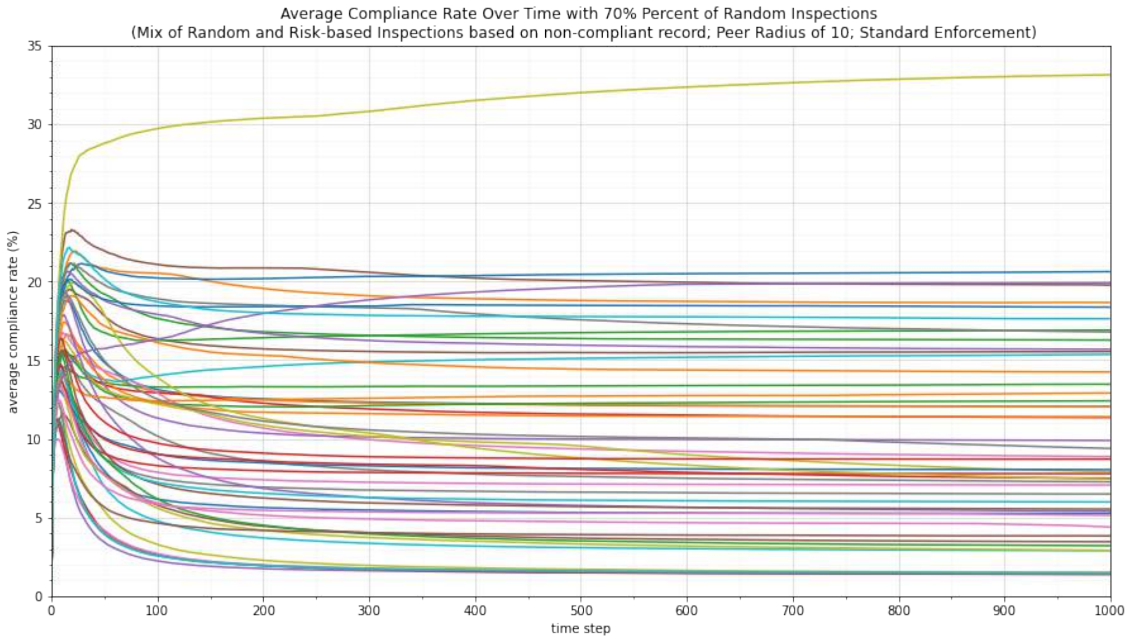

(g) 70% random inspections (SE).

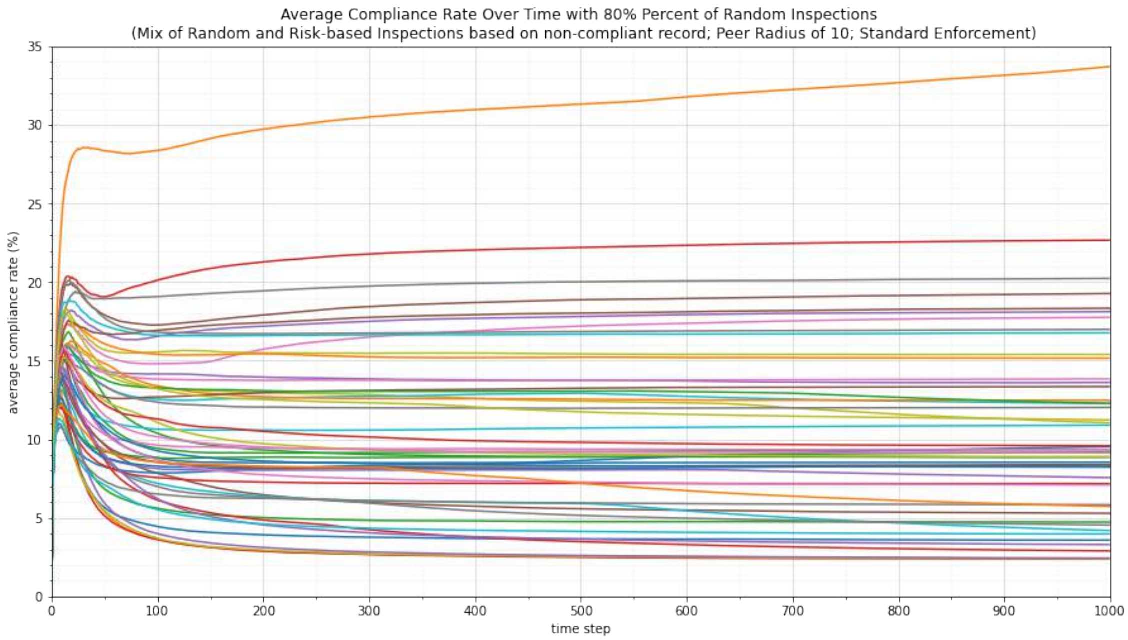

(h) 80% random inspections (SE).

Figure A.17: Average compliance rate results with varying $\%_{rand}$ (SE, MRRbNC, continued).

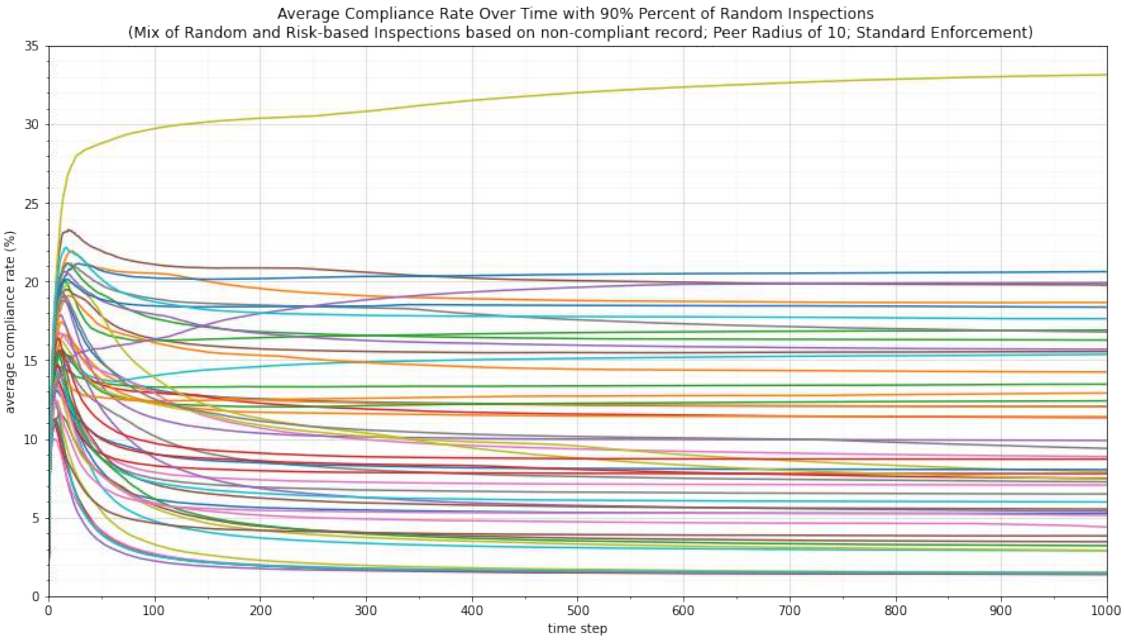

(i) 90% random inspections (SE).

Figure A.17: Average compliance rate results with varying $\%_{rand}$ (SE, MRRbNC, continued).

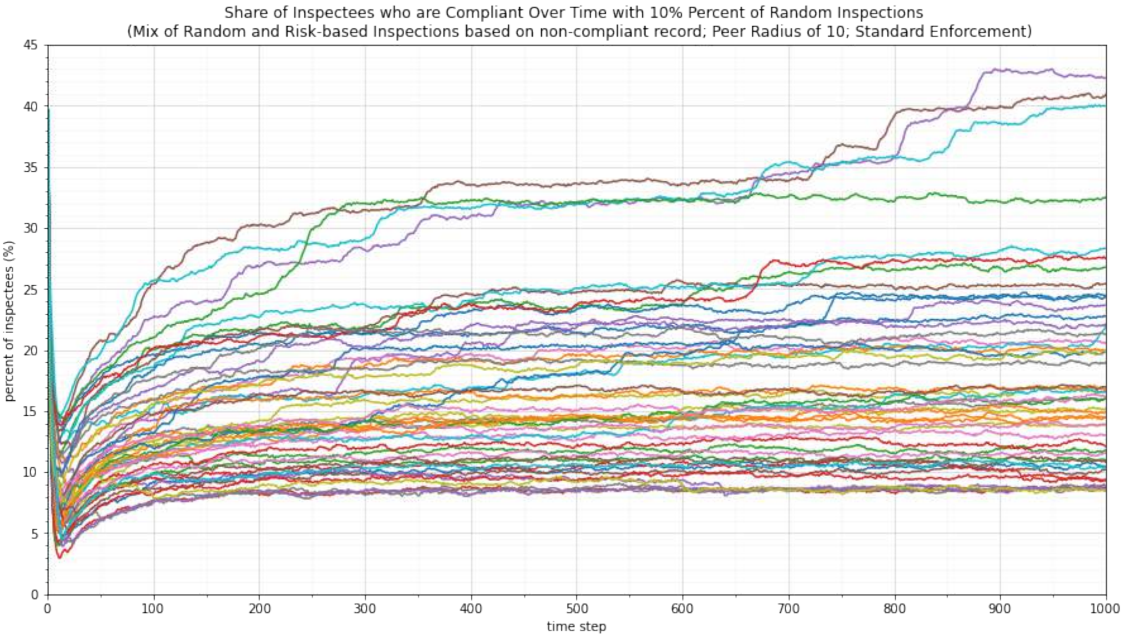

(a) 10% random inspections (SE).

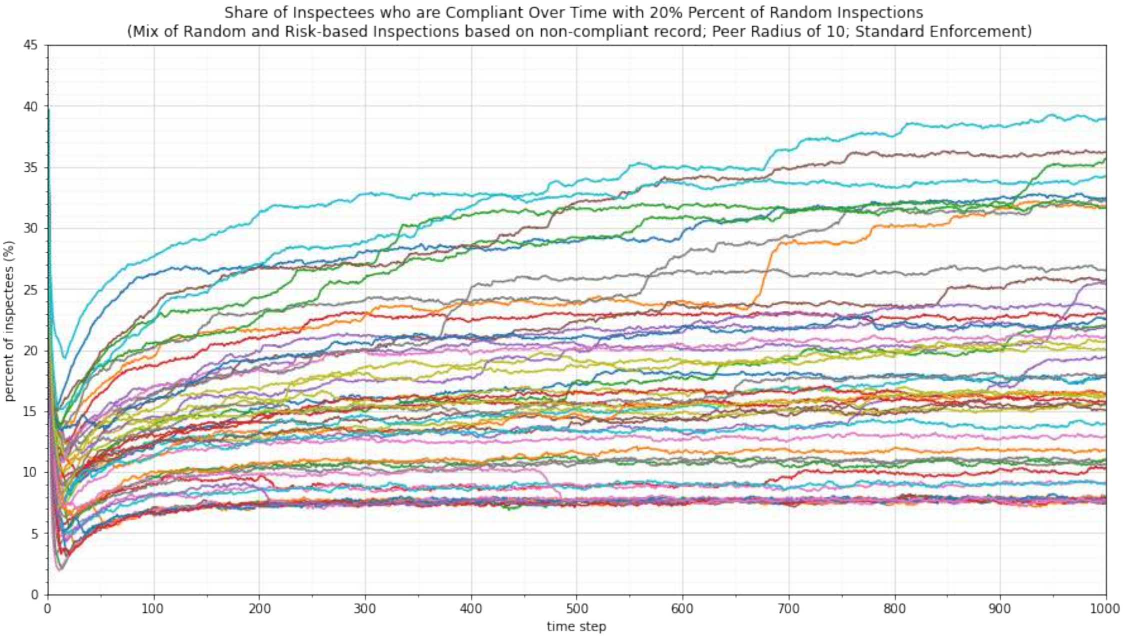

(b) 20% random inspections (SE).

Figure A.18: Share of compliant inspectees with varying $\%_{rand}$ (SE, MRRbNC).

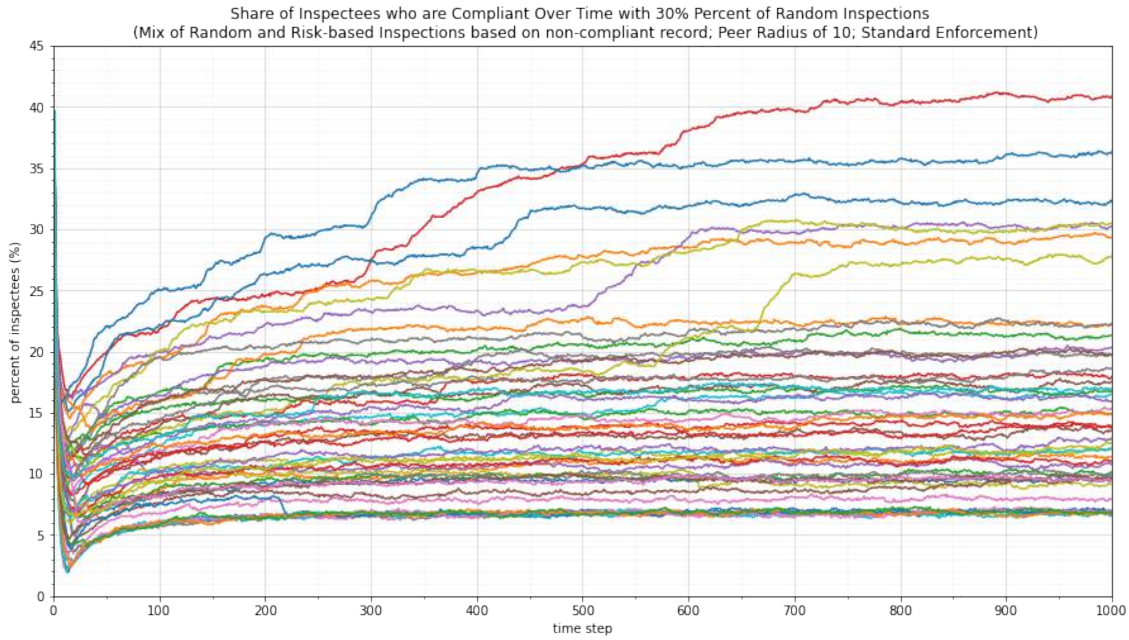

(c) 30% random inspections (SE).

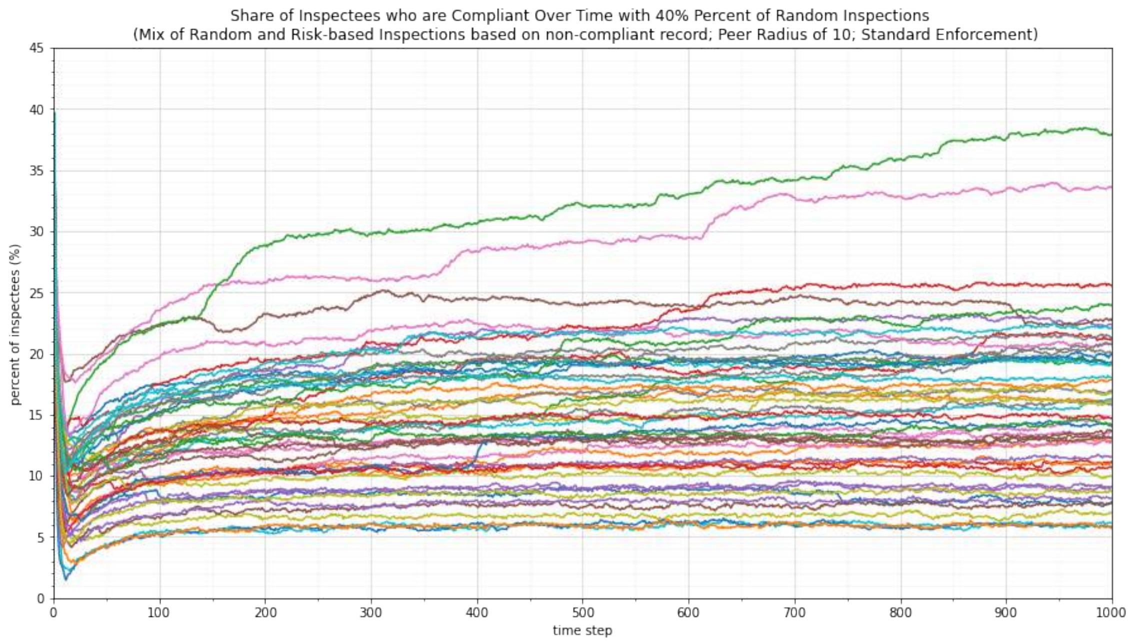

(d) 40% random inspections (SE).

Figure A.18: Share of compliant inspectees with varying p$\%_{rand}$ (SE, MRRbNC, continued).

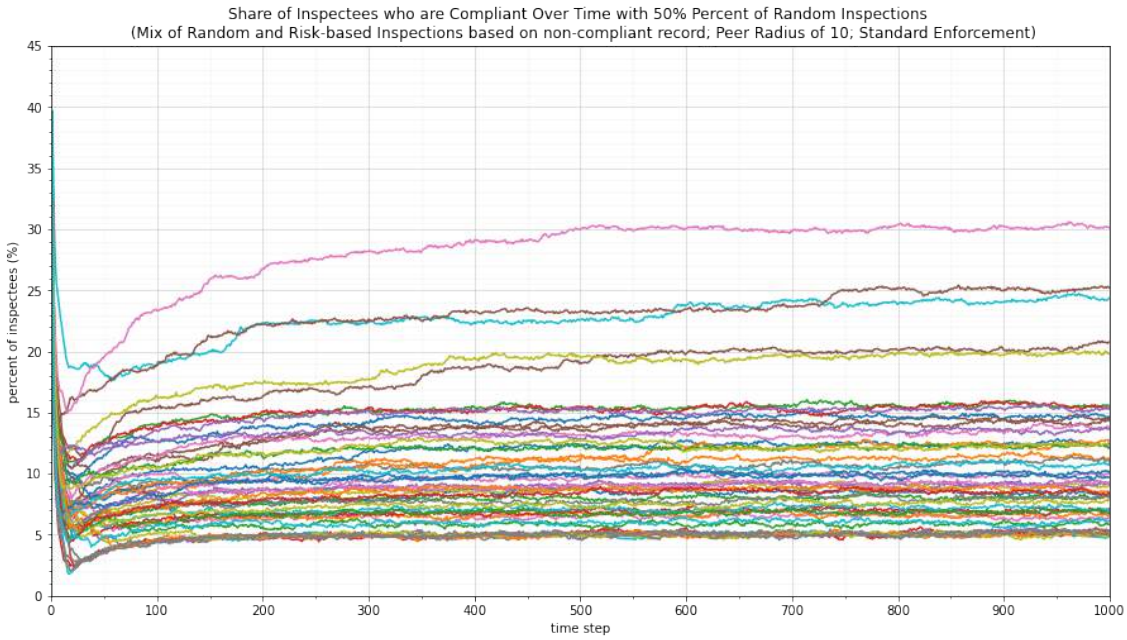

(e) 50% random inspections (SE).

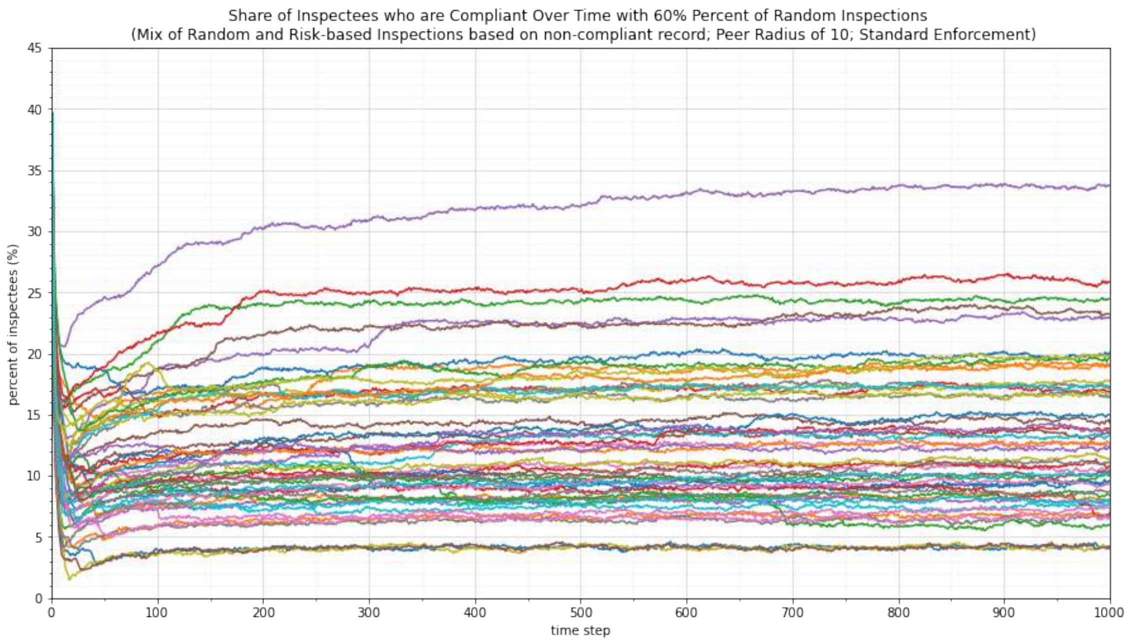

(f) 60% random inspections (SE).

Figure A.18: Share of compliant inspectees with varying $\%_{rand}$ (SE, MRRbNC, continued).

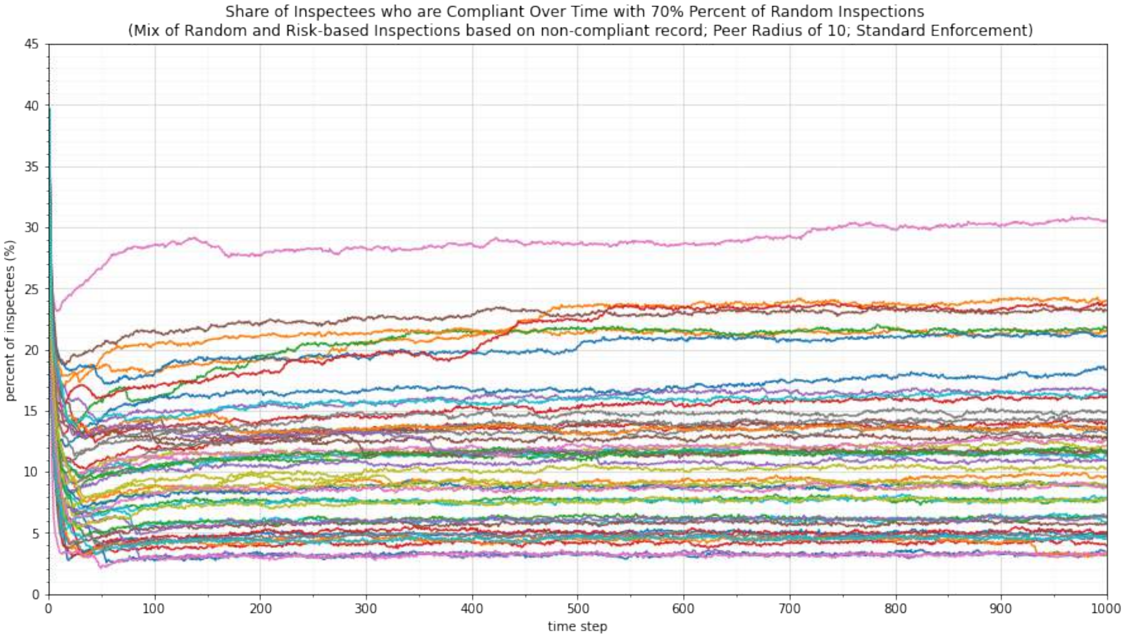

(g) 70% random inspections (SE).

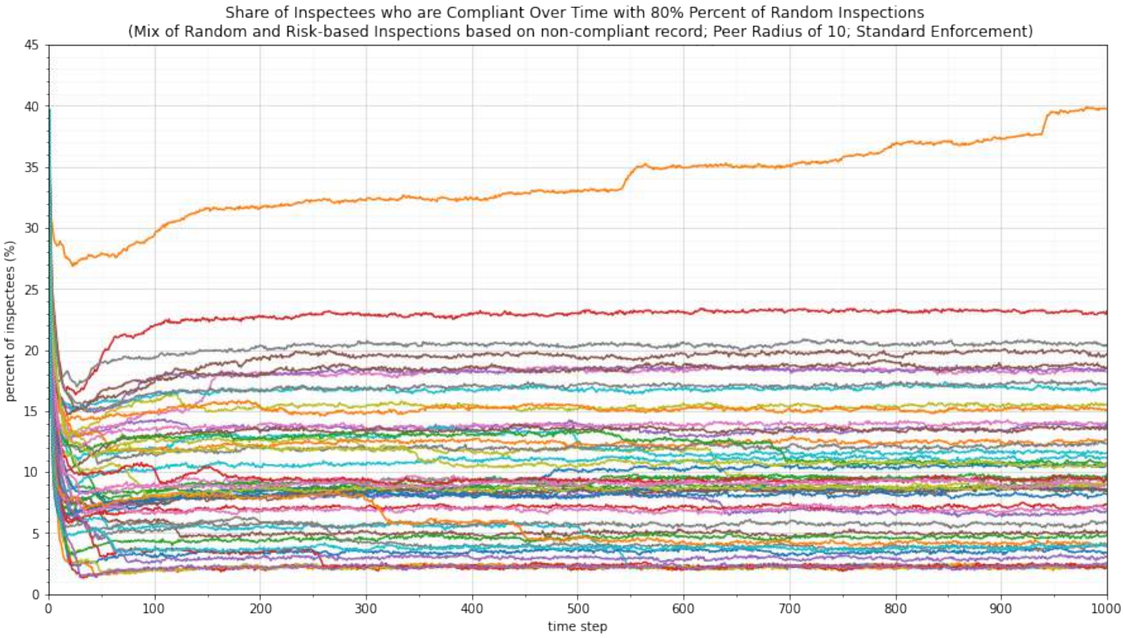

(h) 80% random inspections (SE).

Figure A.18: Share of compliant inspectees with varying $\%_{rand}$ (SE, MRRbNC, continued).

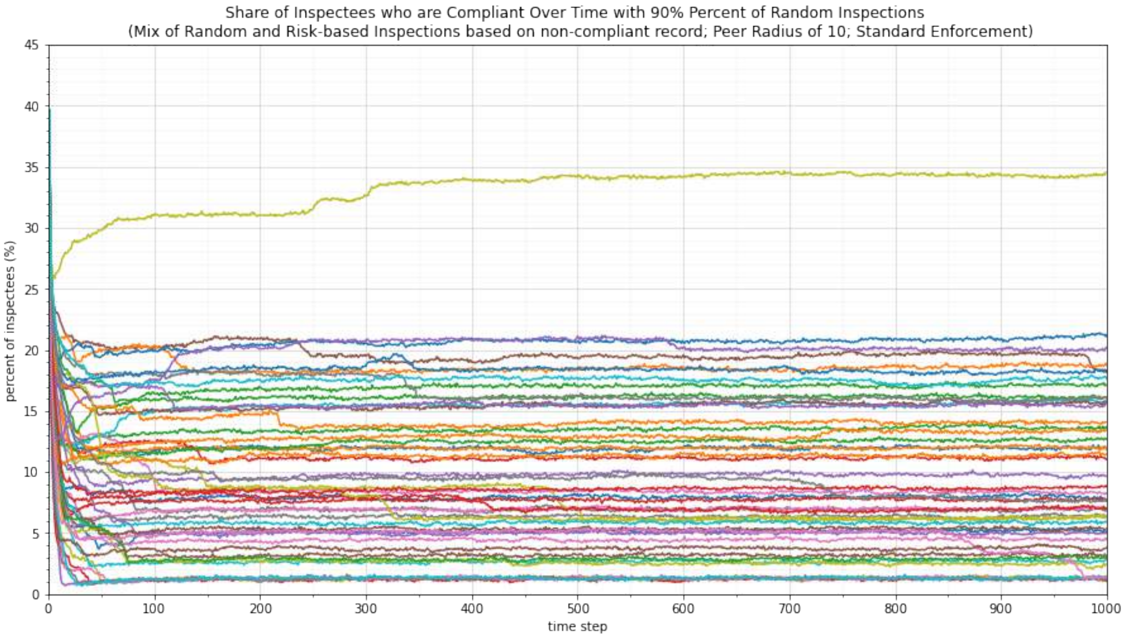

(i) 90% random inspections (SE).

Figure A.18: Share of compliant inspectees with varying $\%_{rand}$ (SE, MRRbNC, continued).

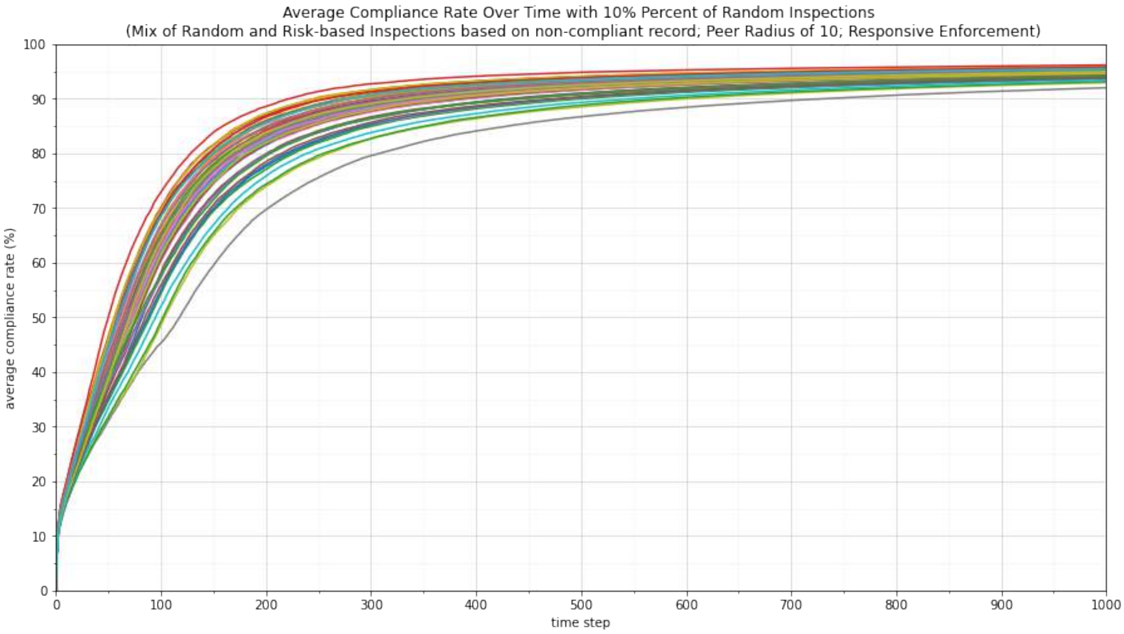

(a) 10% random inspections (RE).

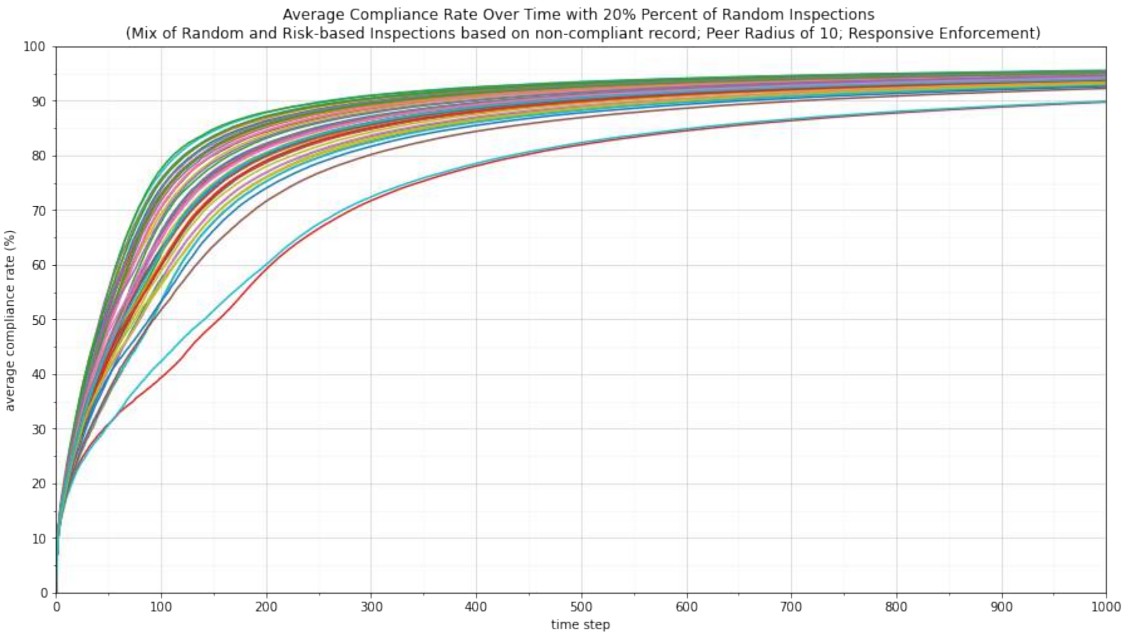

(b) 20% random inspections (RE).

Figure A.19: Average compliance rate results with varying $\%_{rand}$ (RE).

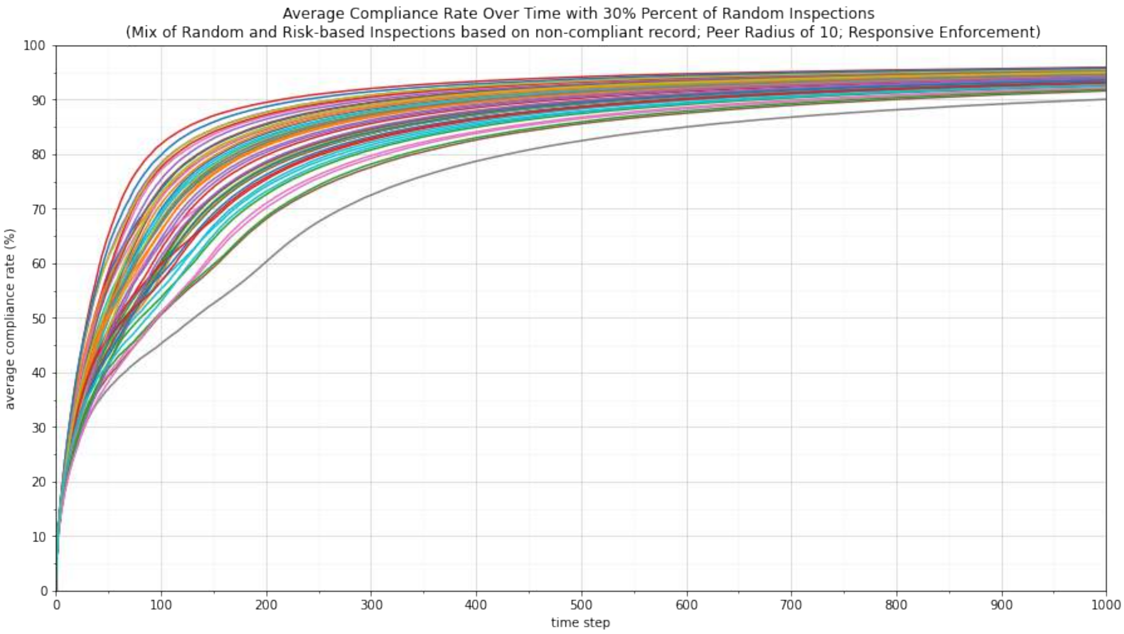

(c) 30% random inspections (RE).

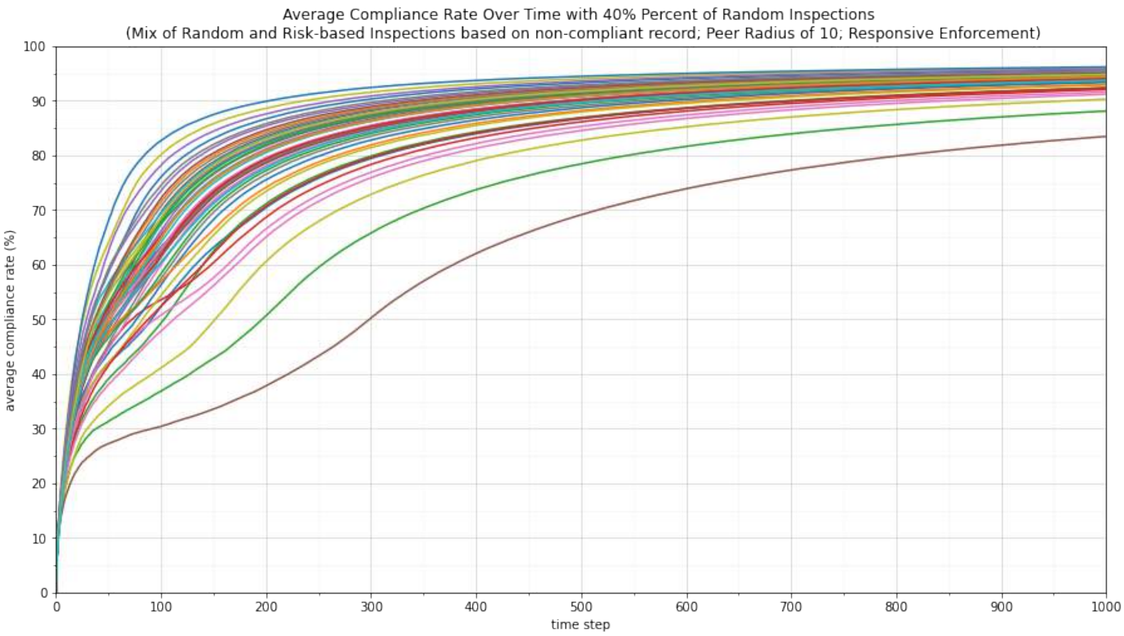

(d) 40% random inspections (RE).

Figure A.19: Average compliance rate results with varying $\%_{rand}$ (RE) (continued).

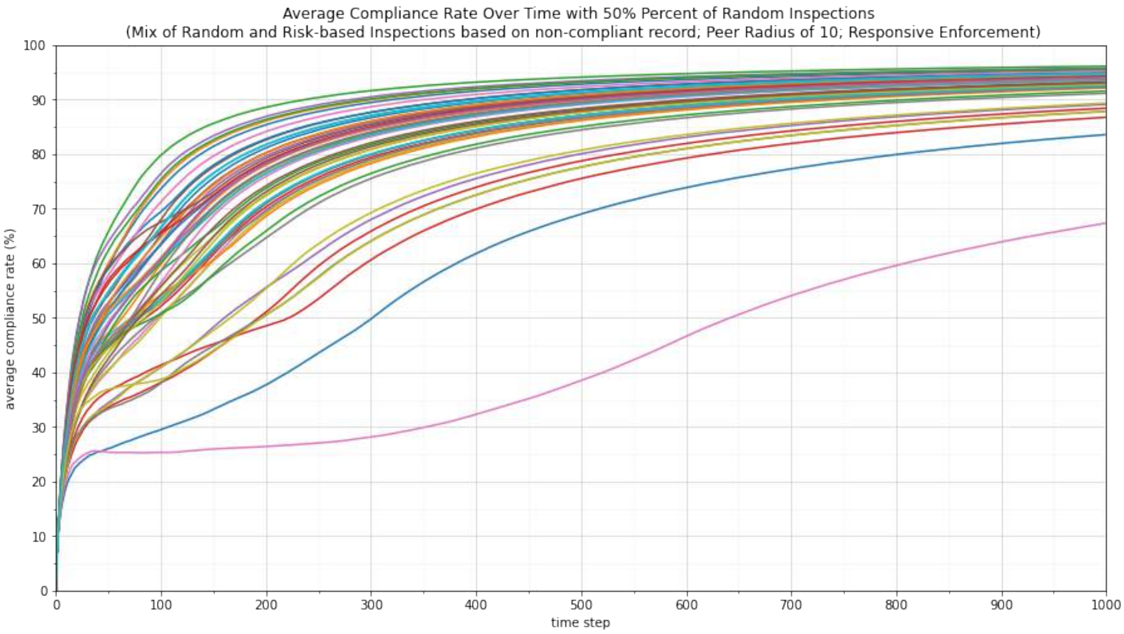

(e) 50% random inspections (RE).

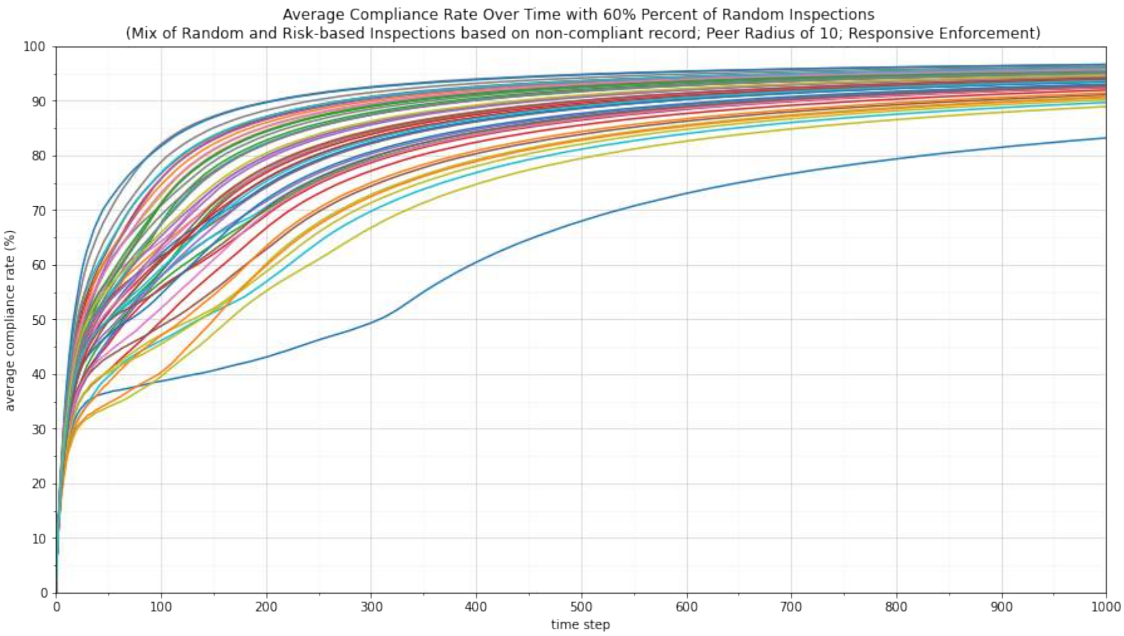

(f) 60% random inspections (RE).

Figure A.19: Average compliance rate results with varying $\%_{rand}$ (RE) (continued).

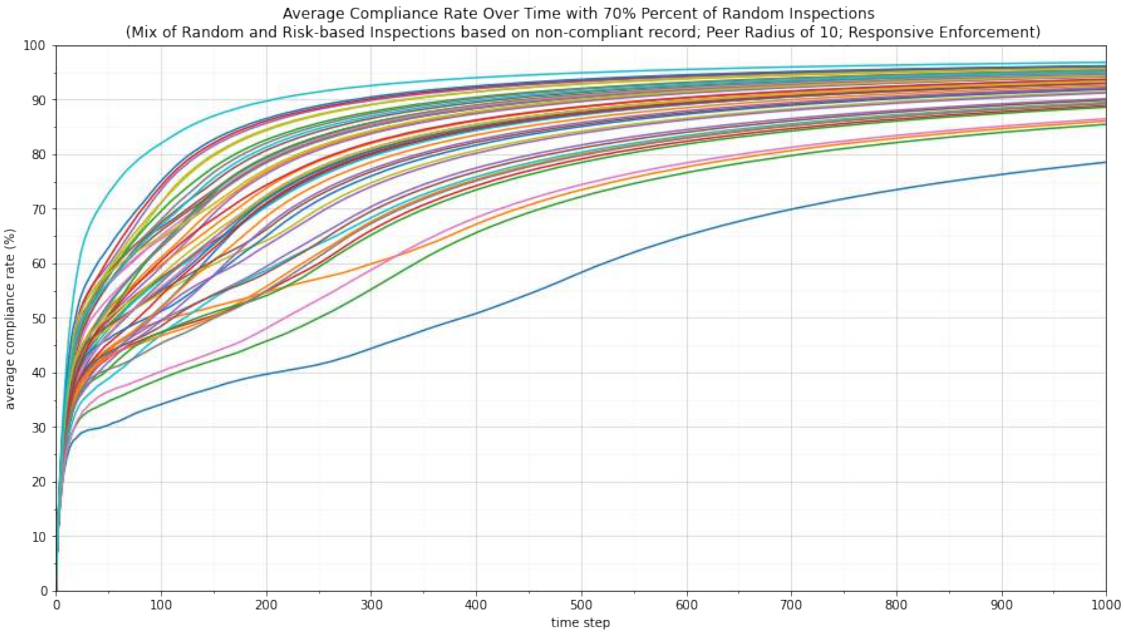

(g) 70% random inspections (RE).

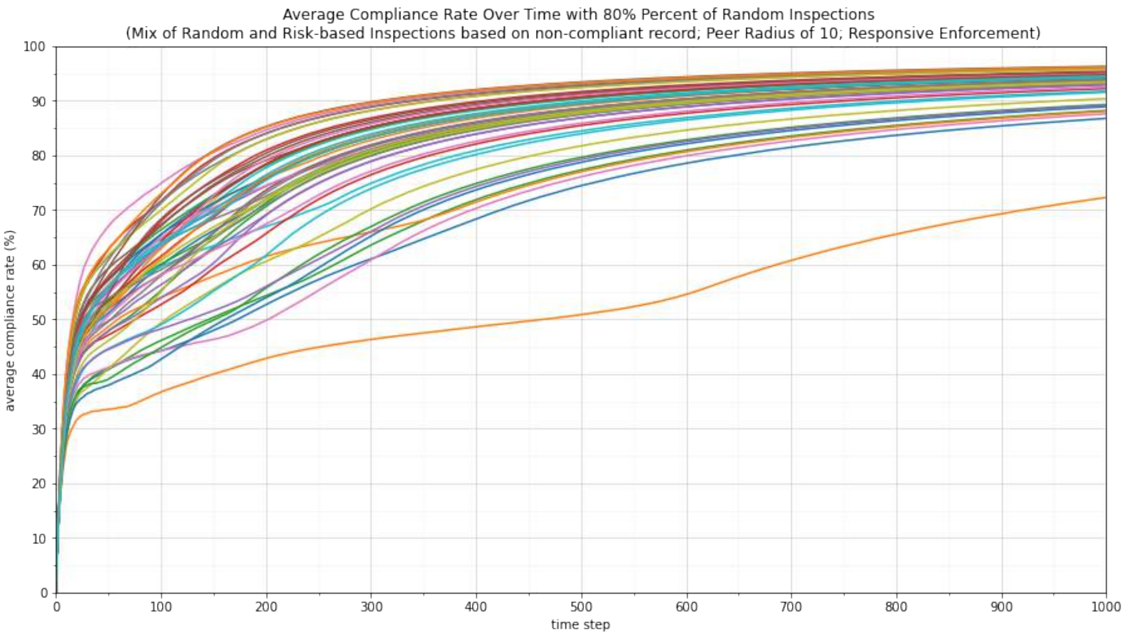

(h) 80% random inspections (RE).

Figure A.19: Average compliance rate results with varying $\%_{rand}$ (RE) (continued).

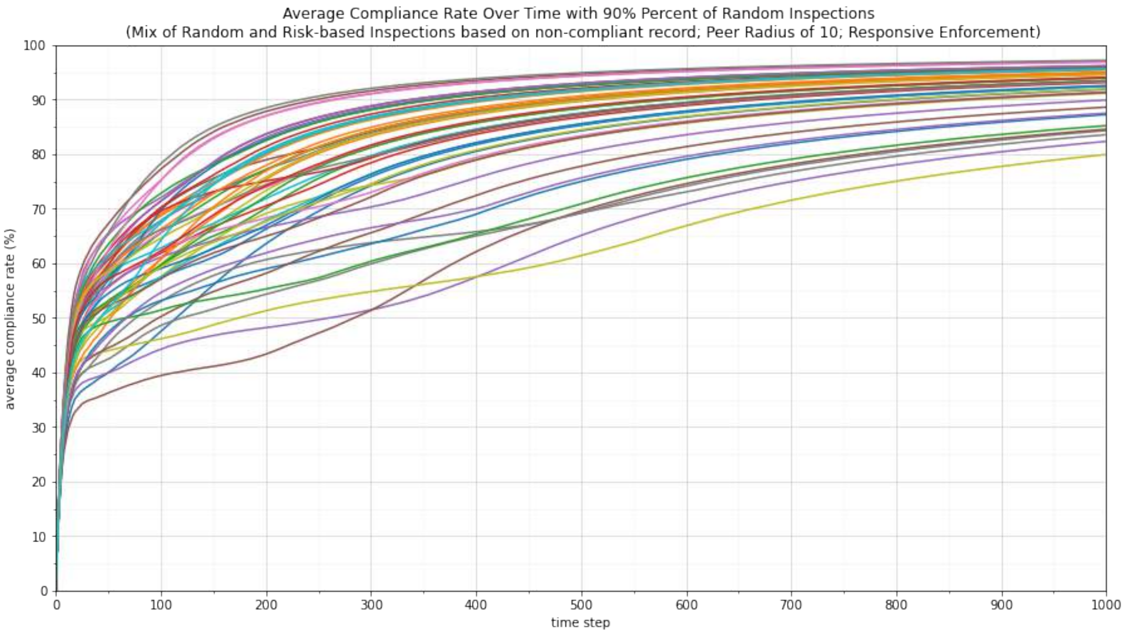

(i) 90% random inspections (RE).

Figure A.19: Average compliance rate results with varying $\%_{rand}$ (RE) (continued).

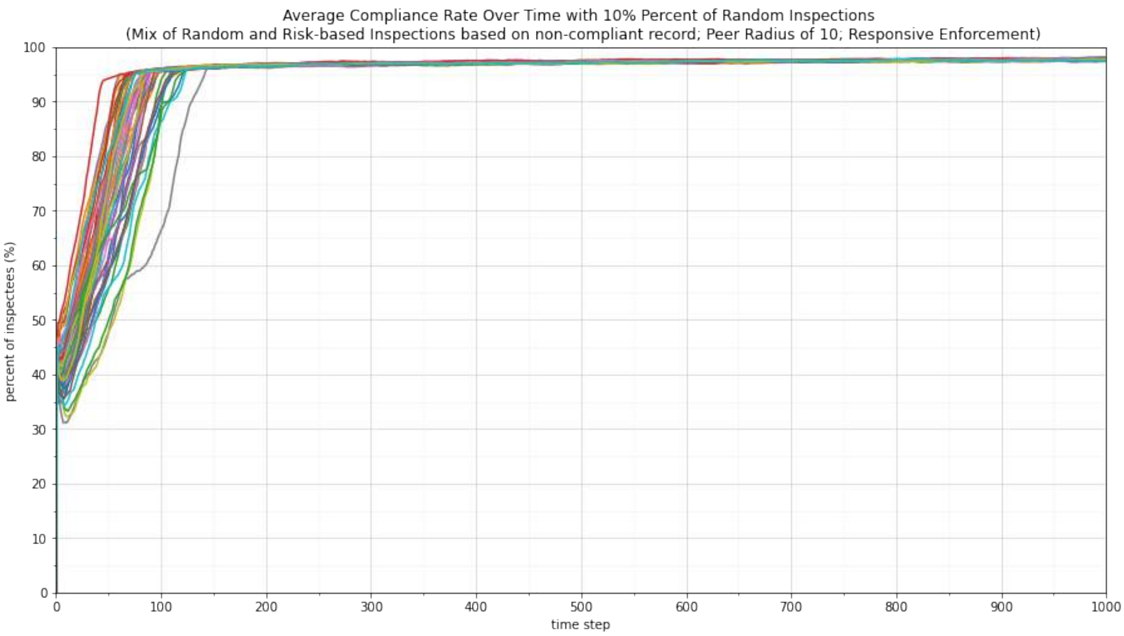

(a) 10% random inspections (RE).

Figure A.20: Share of compliant inspectees with varying $\%_{rand}$ (RE, MRRbNC).

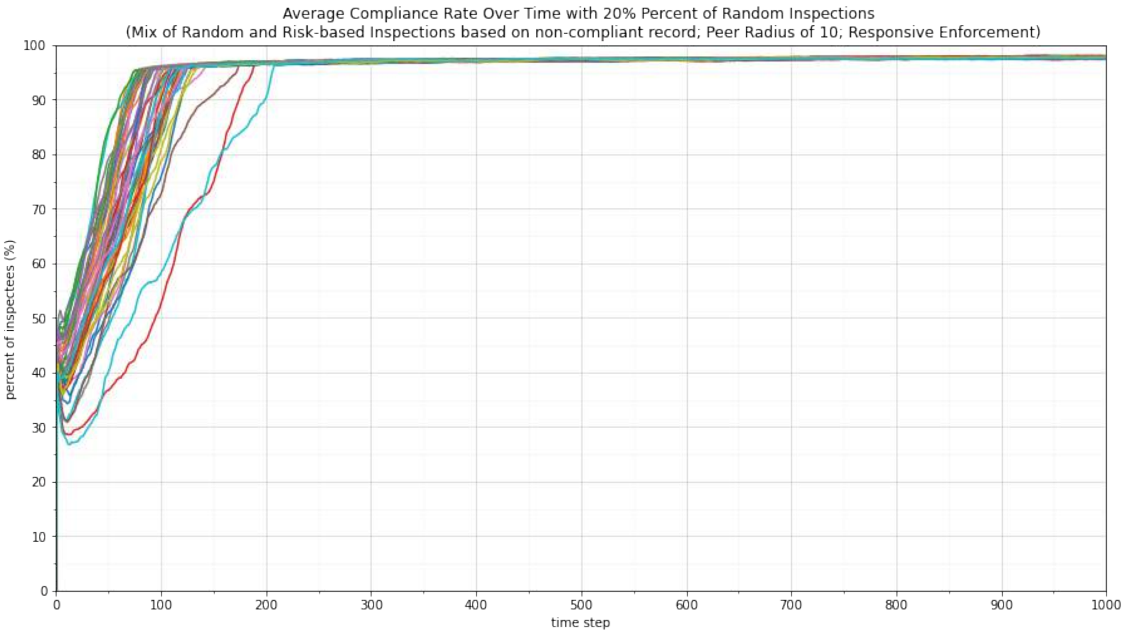

(b) 20% random inspections (RE).

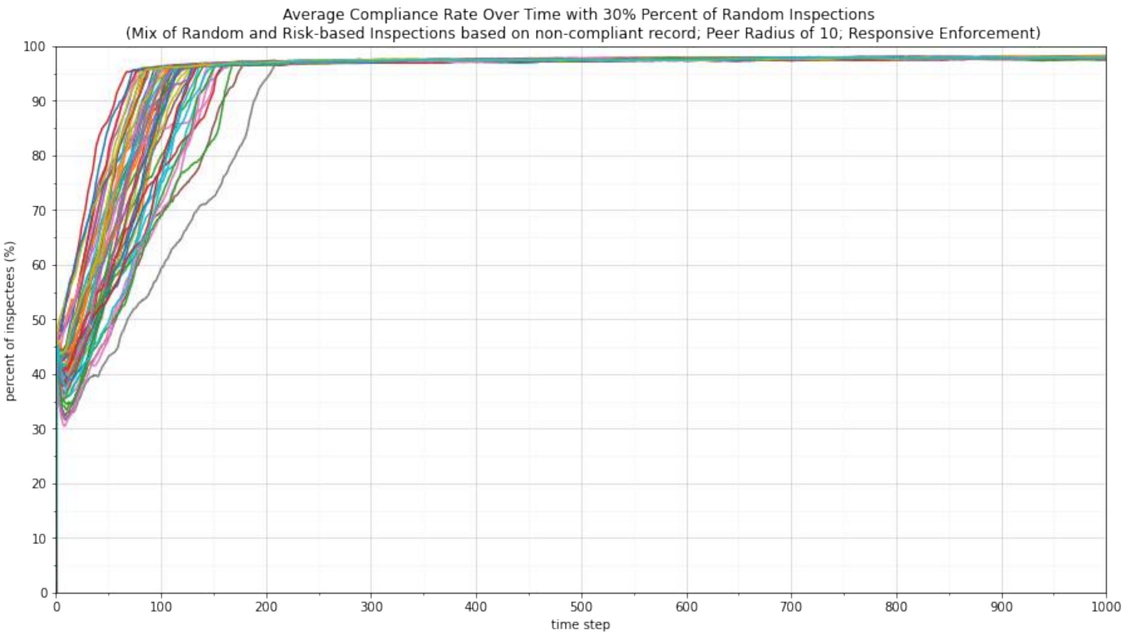

(c) 30% random inspections (RE).

Figure A.20: Share of compliant inspectees with varying $\%_{rand}$ (RE, MRRbNC, continued).

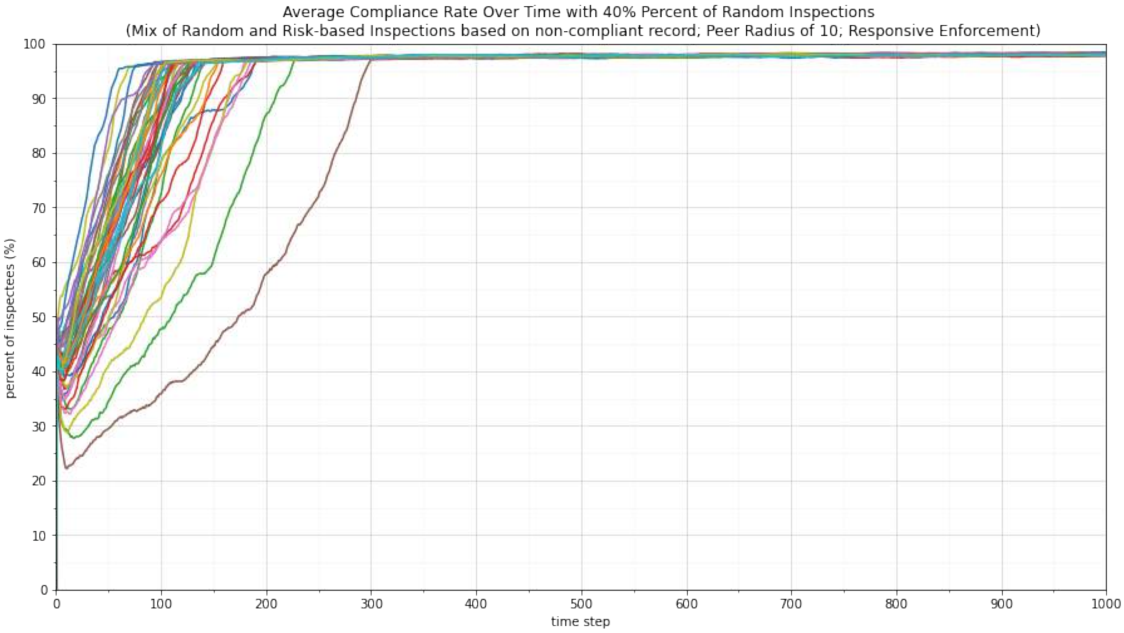

(d) 40% random inspections (RE).

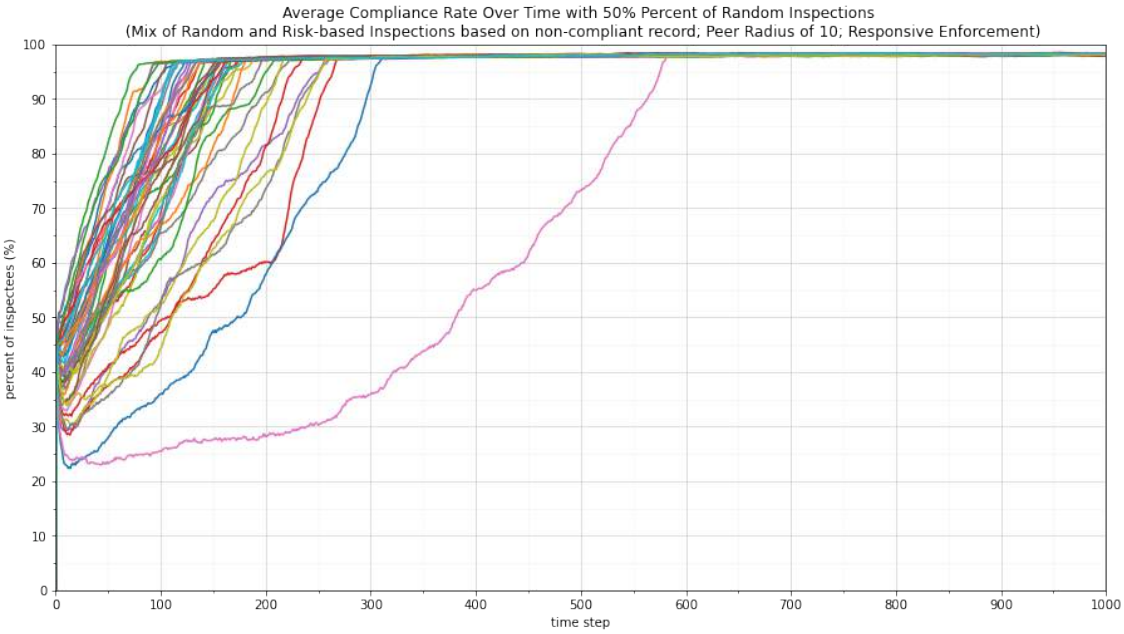

(e) 50% random inspections (RE).

Figure A.20: Share of compliant inspectees with varying $\%_{rand}$ (RE, MRRbNC, continued).

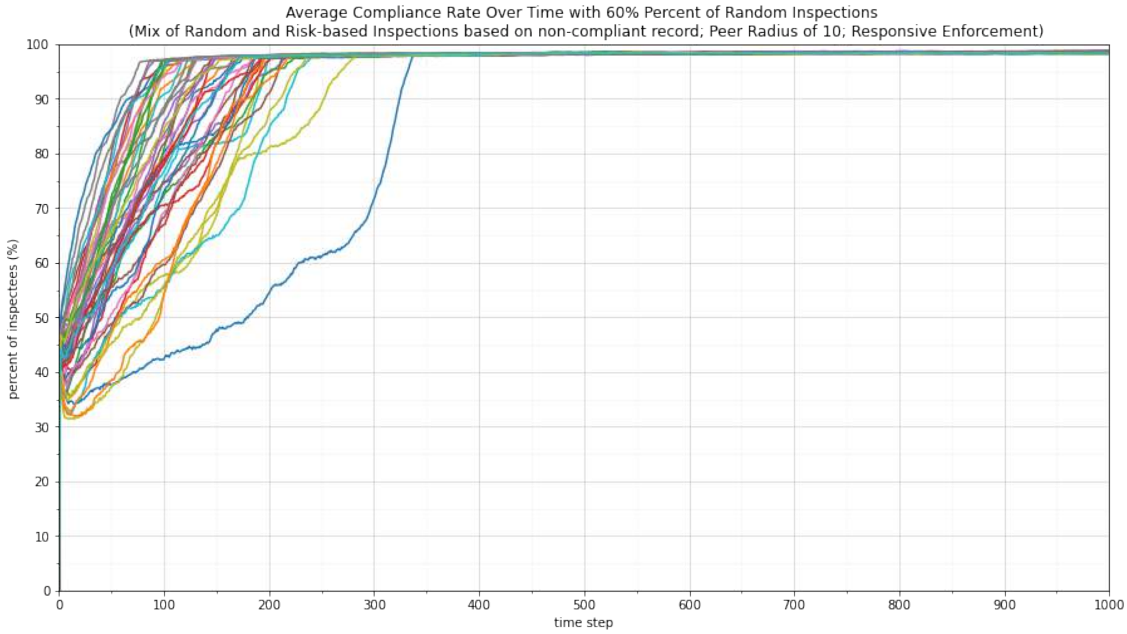

(f) 60% random inspections (RE).

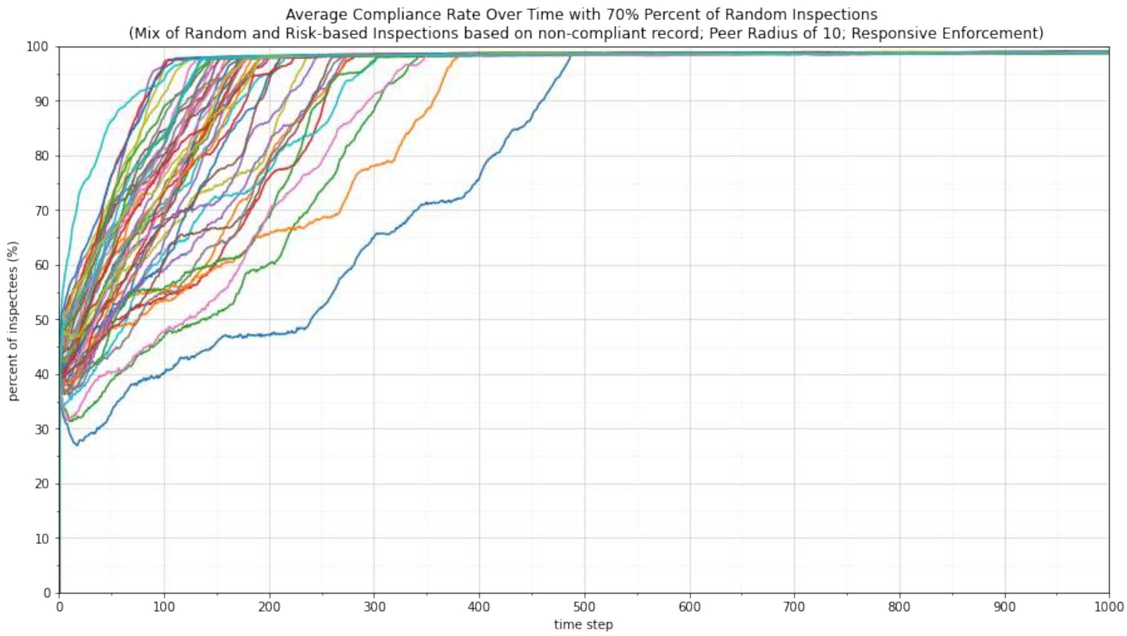

(g) 70% random inspections (RE).

Figure A.20: Share of compliant inspectees with varying $\%_{rand}$ (RE, MRRbNC, continued).

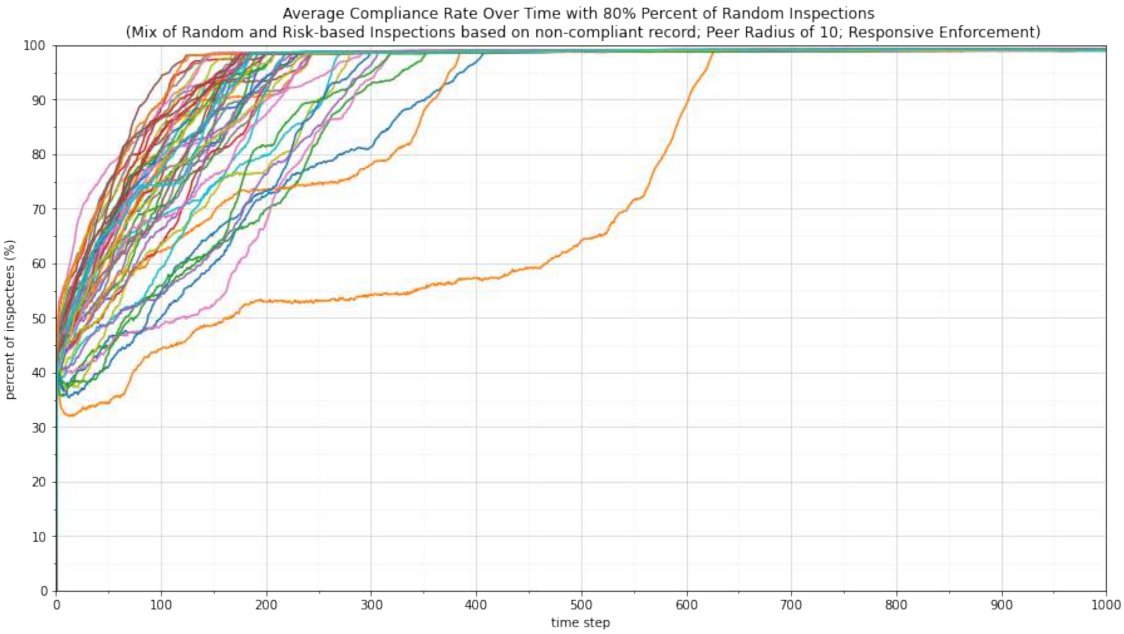

(h) 80% random inspections (RE).

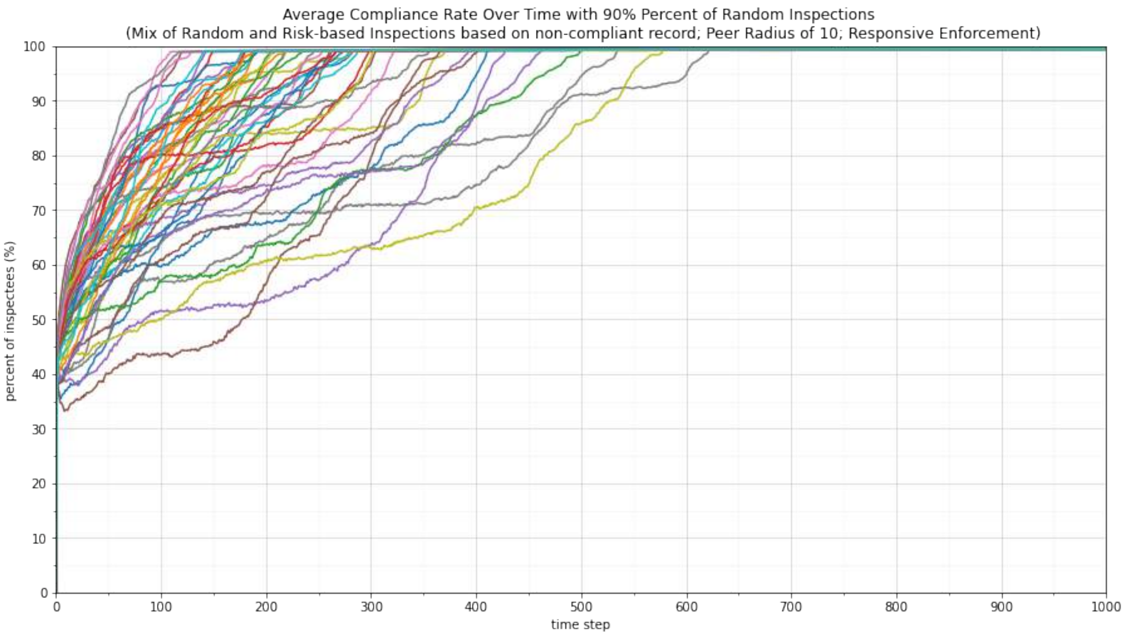

(i) 90% random inspections (RE).

Figure A.20: Share of compliant inspectees with varying $\%_{rand}$ (RE, MRRbNC, continued).

A.1.2.2 Mixed Random and Risk-based Inspection Strategy based on Offense (MRRbOS)

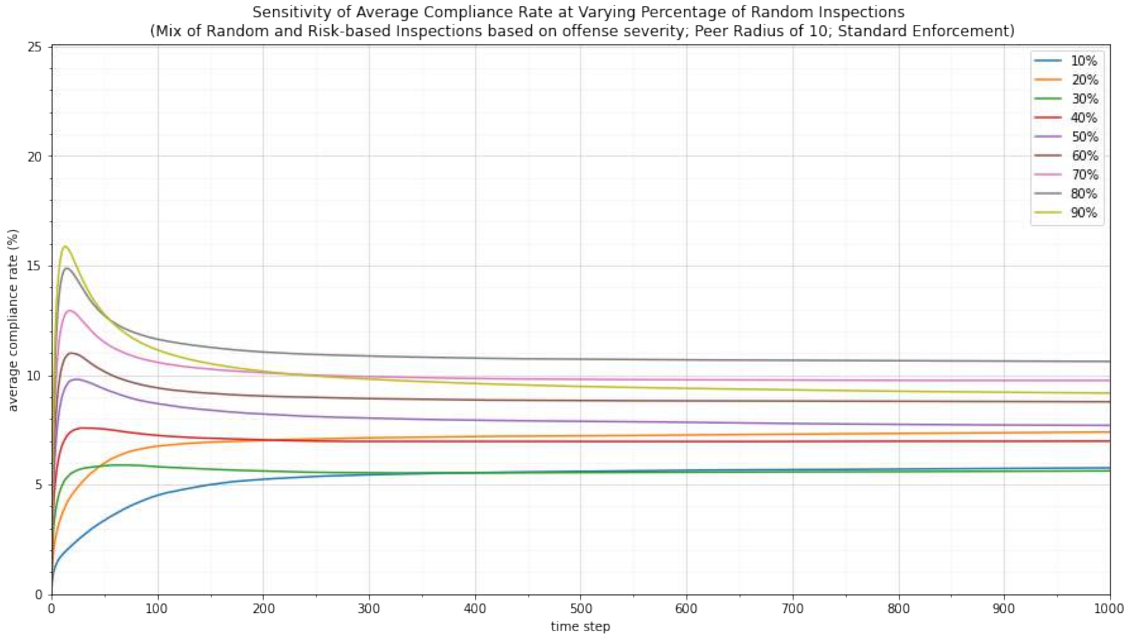

(a) Standard enforcement.

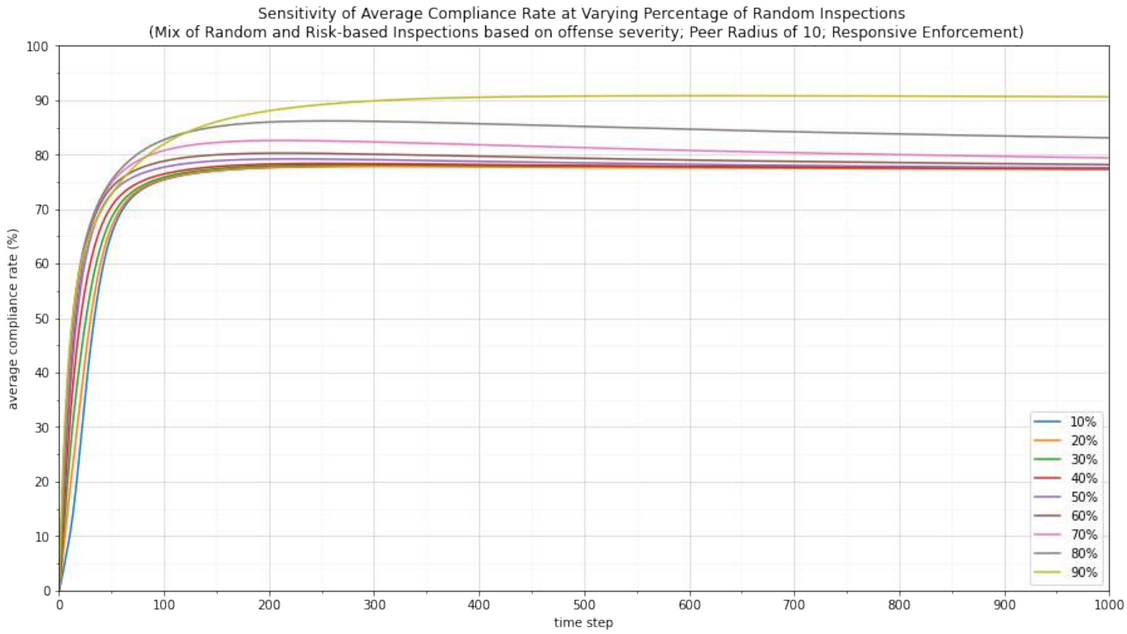

(b) Responsive enforcement.

Figure A.21: Average compliance rate results with varying $\%_{rand}$ (MRRbOS).

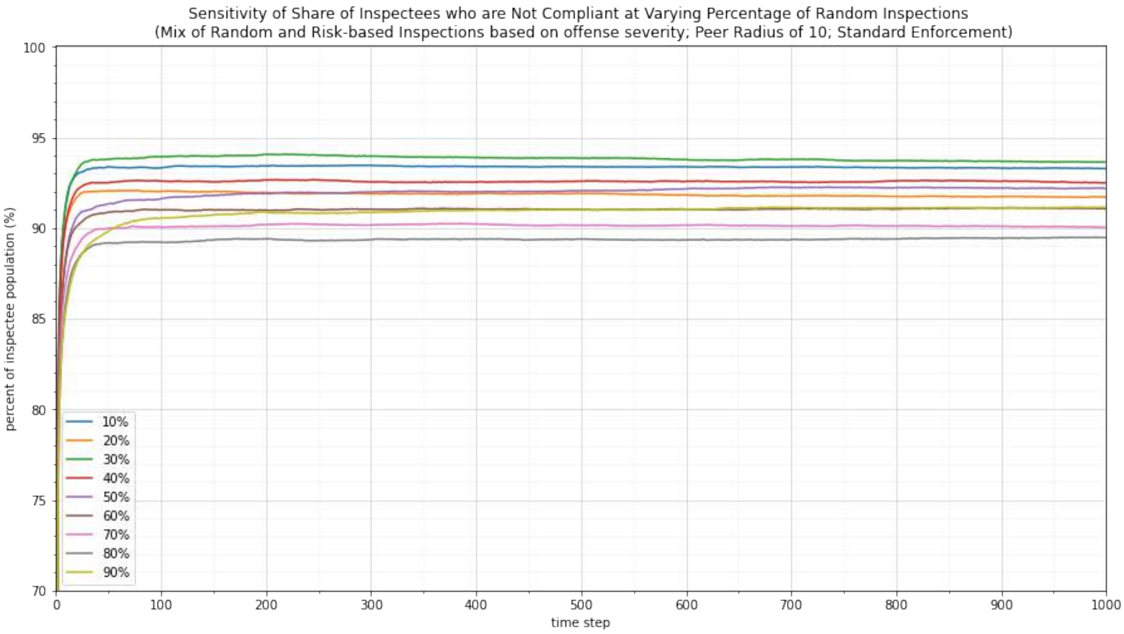

(a) Standard enforcement.
For a detailed breakdown by compliance level, see Figure A.23.

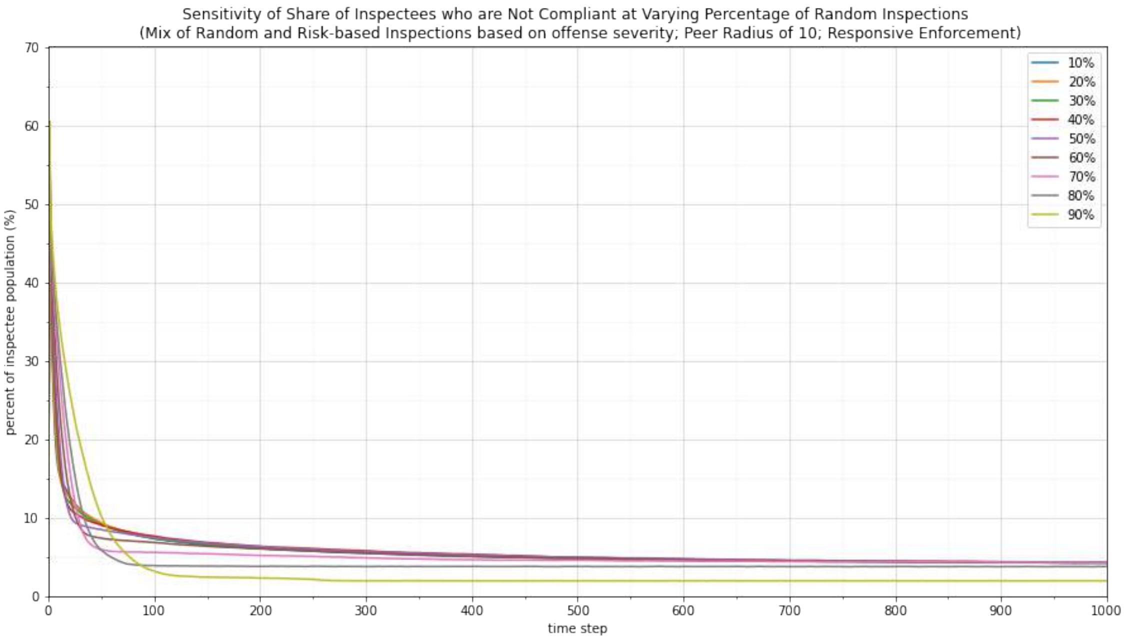

(b) Responsive enforcement.
For a detailed breakdown by compliance level, see Figure A.25.

Figure A.22: Share of non-compliant inspectees with varying $\%_{rand}$ (MRRbOS).

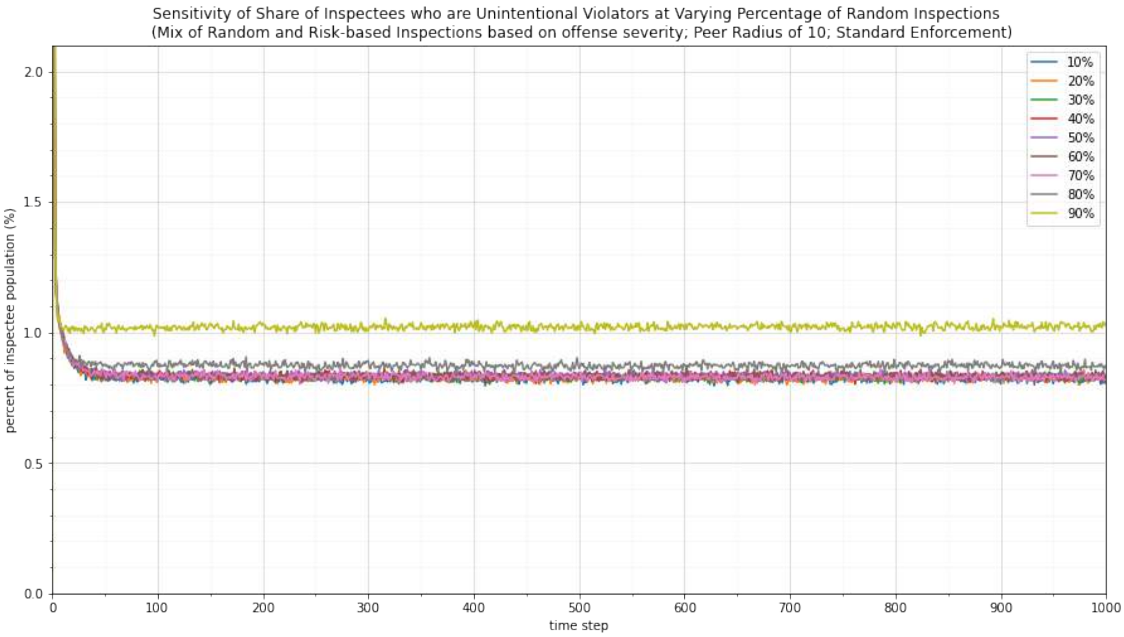

(a) Unintentional violators.

Figure A.23: Non-compliant inspectee breakdown with varying $\%_{rand}$ (SE, MRR-bOS).

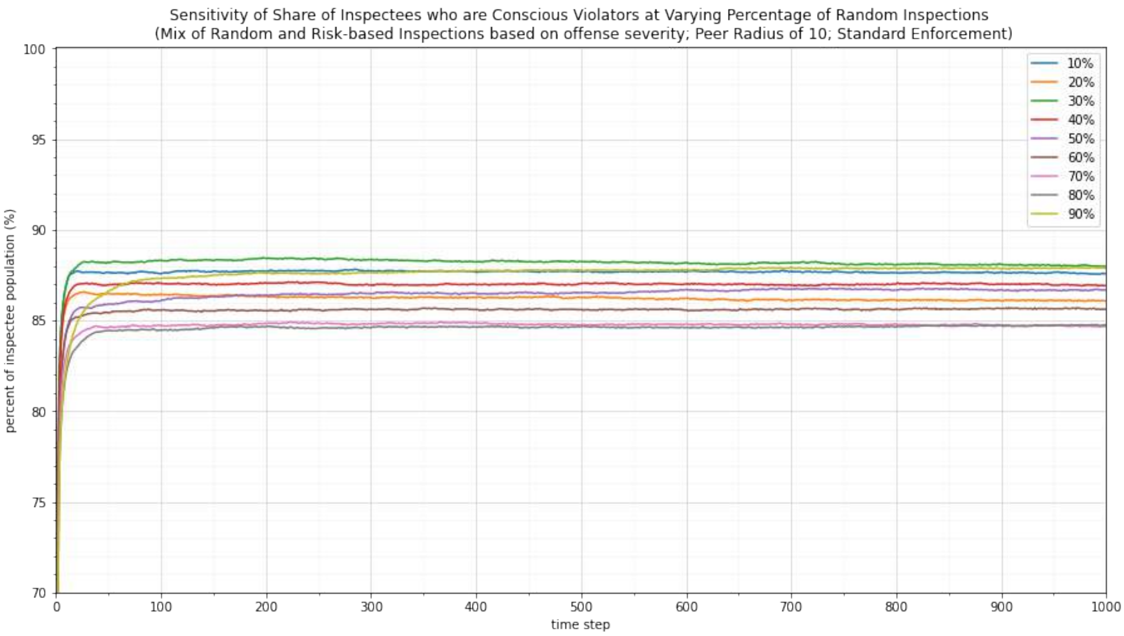

(a) Conscious violators.

Figure A.24: Non-compliant inspectee breakdown with varying $\%_{rand}$ (SE, MRR-bOS, continued).

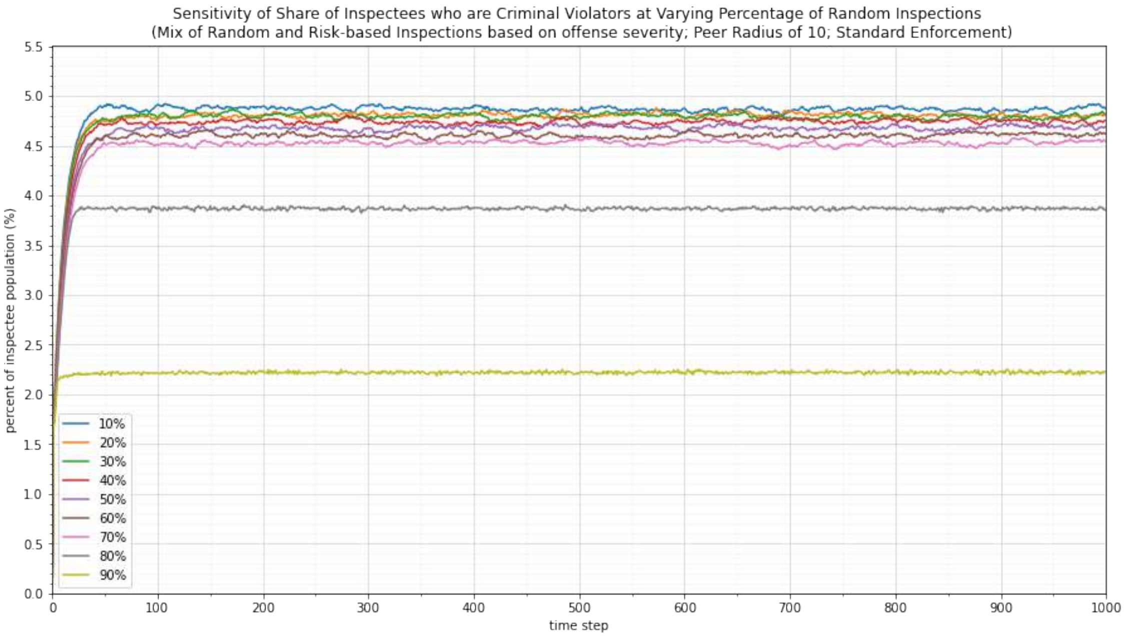

(b) Criminal violators.

Figure A.24: Non-compliant inspectee breakdown with varying $\%_{rand}$ (SE, MRR-bOS, continued).

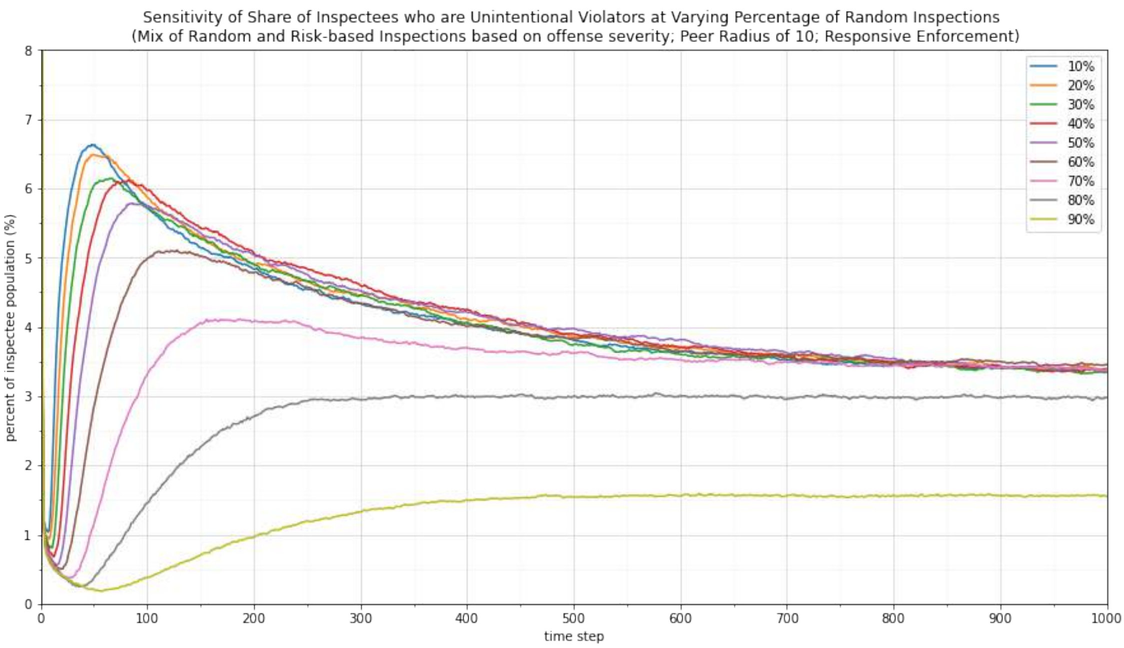

(a) Unintentional violators.

Figure A.25: Non-compliant inspectee breakdown with varying $\%_{rand}$ (RE, MRR-bOS).

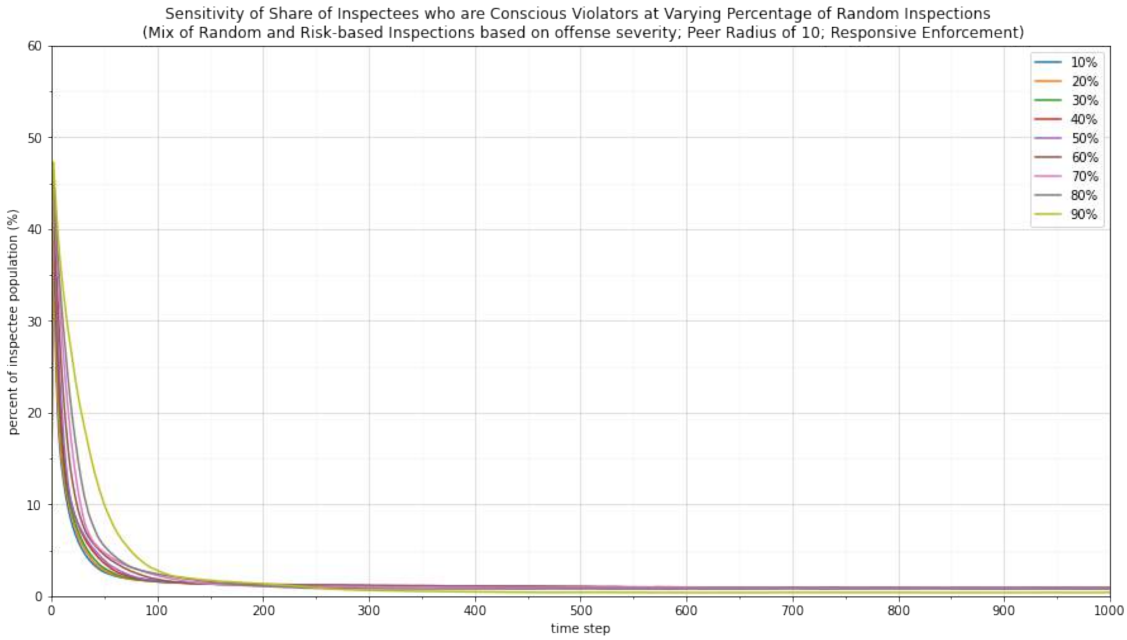

(b) Conscious violators.

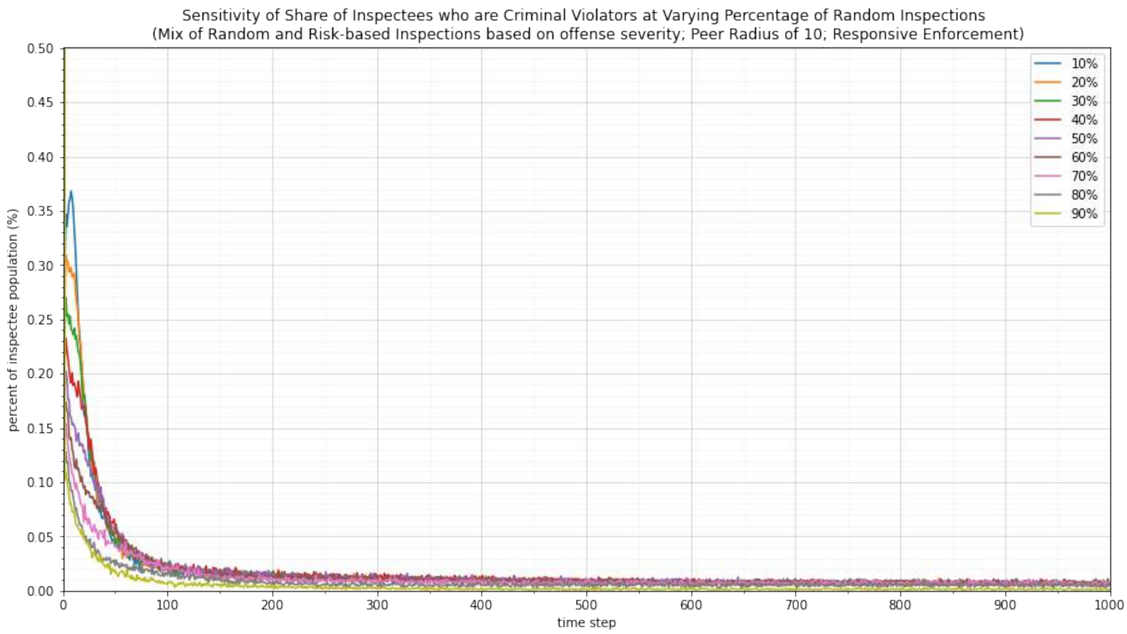

(c) Criminal violators.

Figure A.25: Non-compliant inspectee breakdown with varying $\%_{rand}$ (RE, MRR-bOS, continued).

Path Dependency

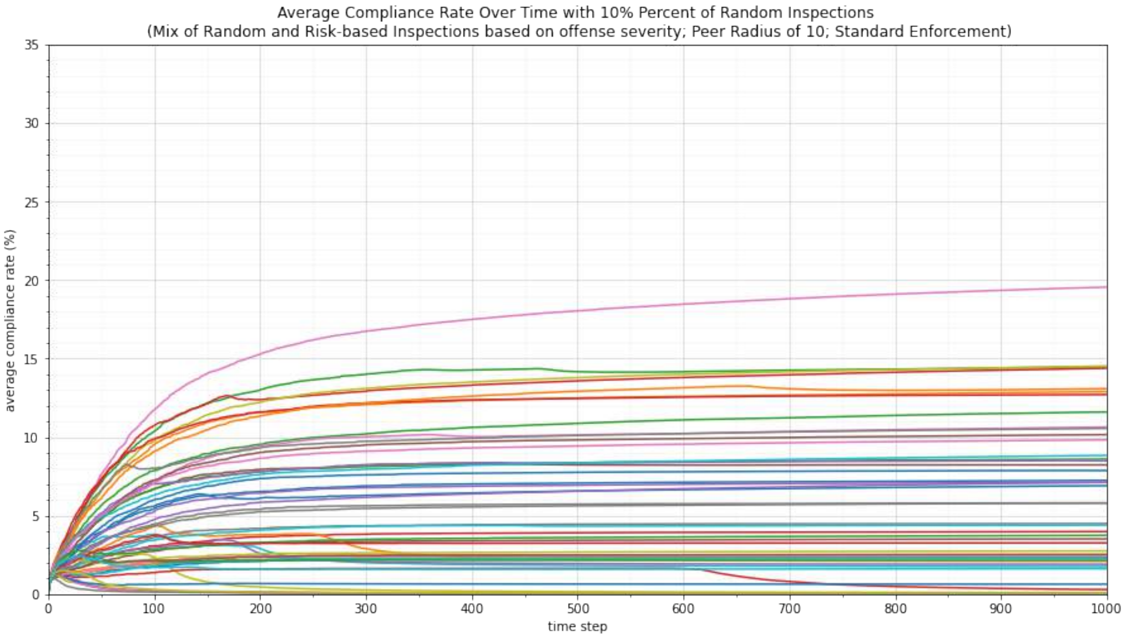

(a) 10% random inspections (SE).

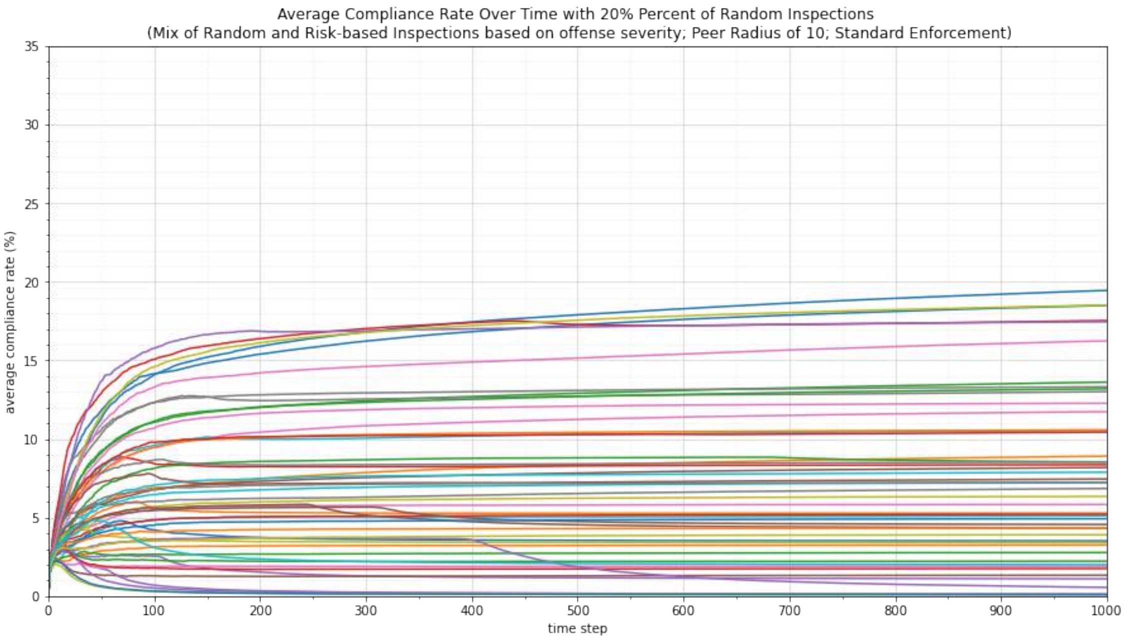

(b) 20% random inspections (SE).

Figure A.26: Average compliance rate results with varying $\%_{rand}$ & SE.

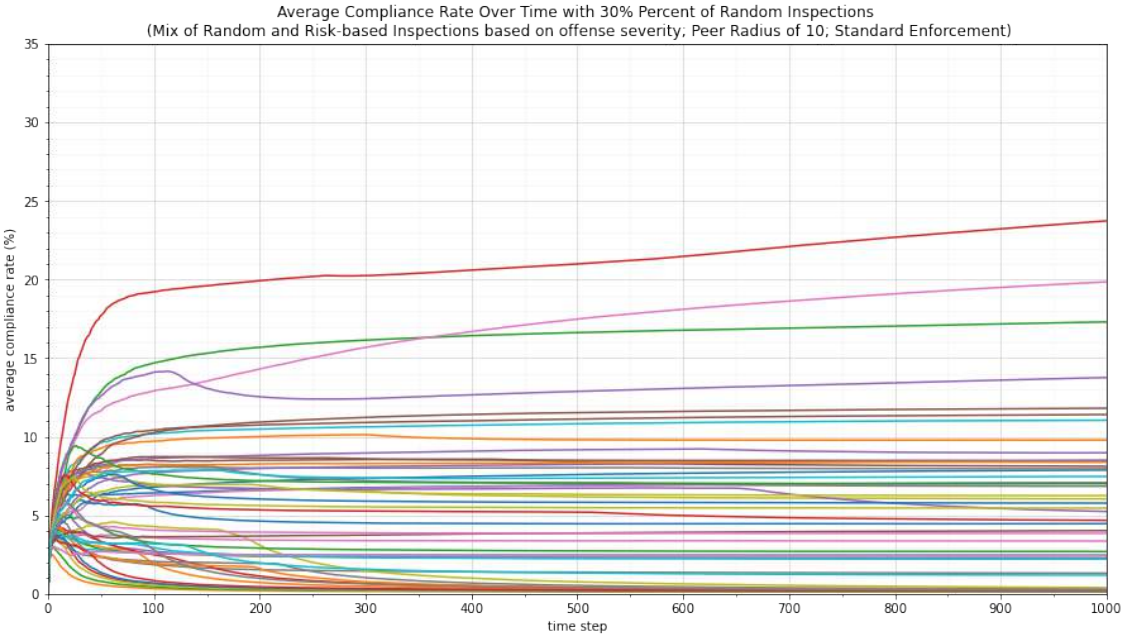

(c) 30% random inspections (SE).

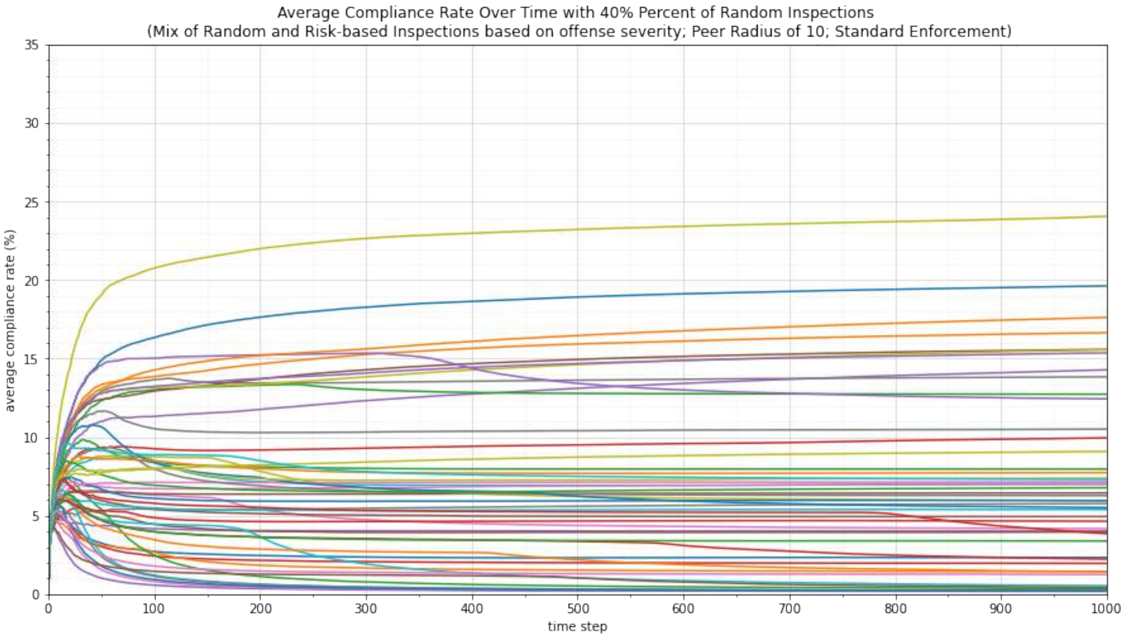

(d) 40% random inspections (SE).

Figure A.26: Average compliance rate results with varying $\%_{rand}$ (SE) (continued).

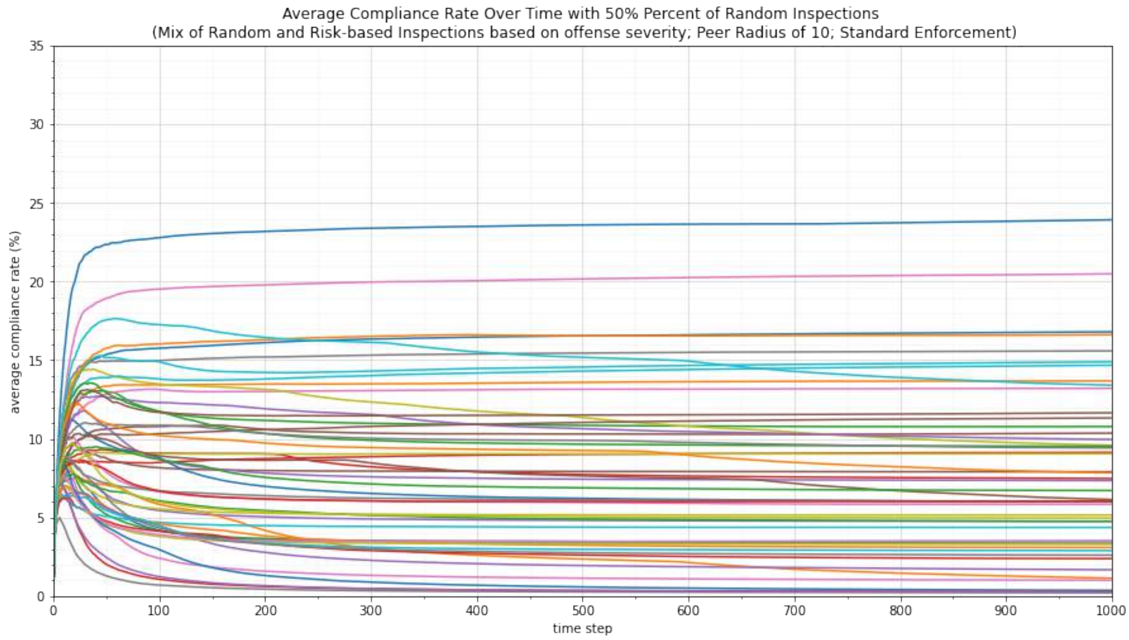

(e) 50% random inspections (SE).

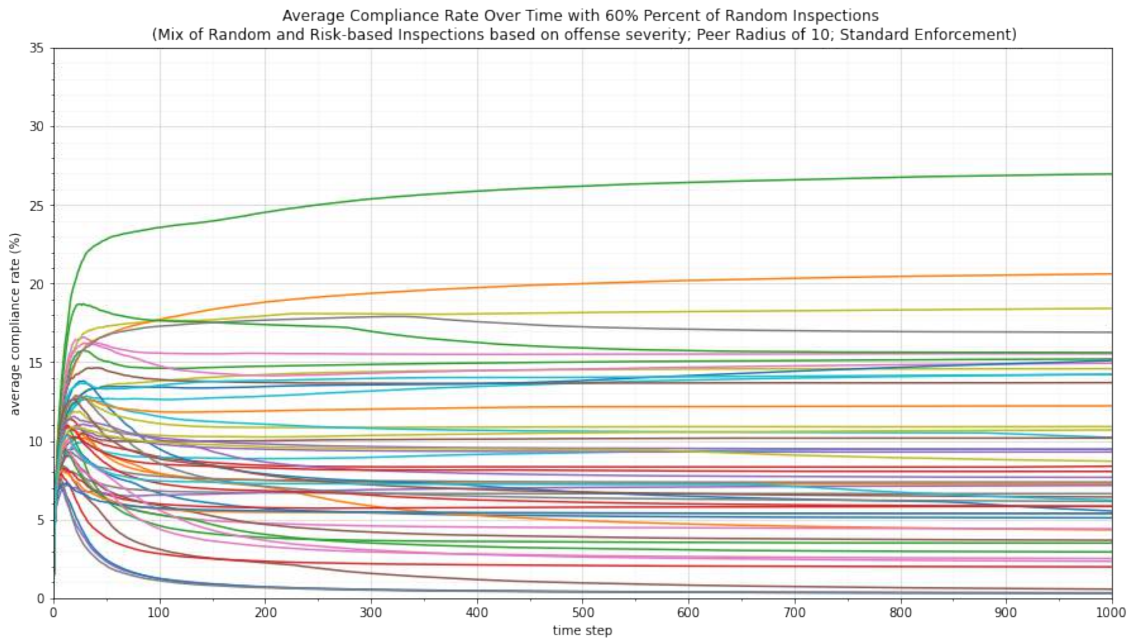

(f) 60% random inspections (SE).

Figure A.26: Average compliance rate results with varying $\%_{rand}$ (SE) (continued).

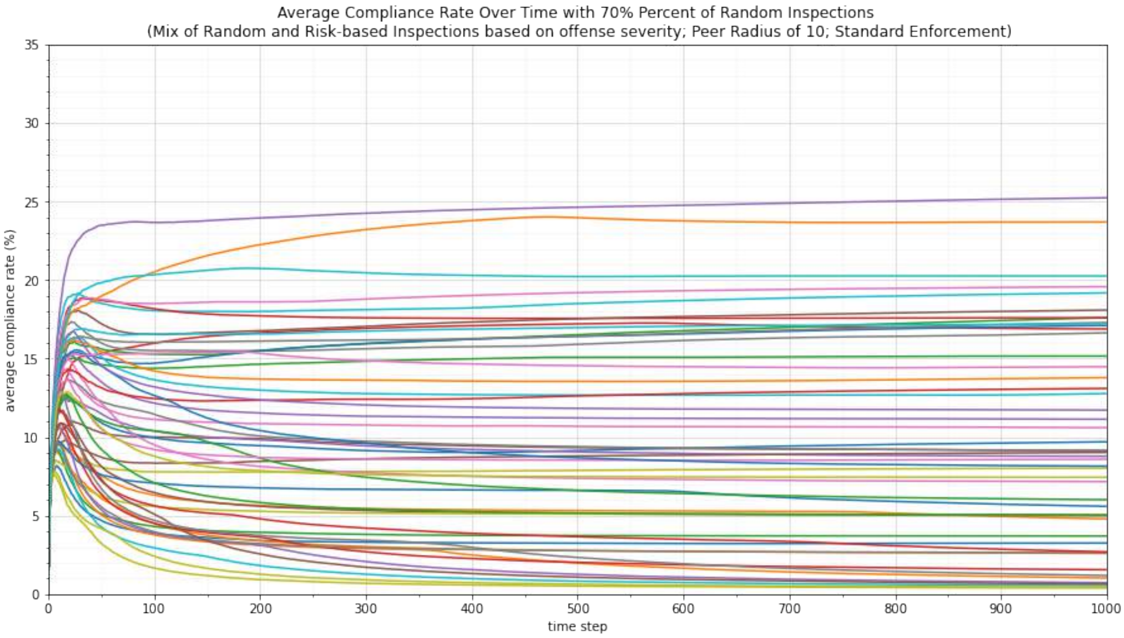

(g) 70% random inspections (SE).

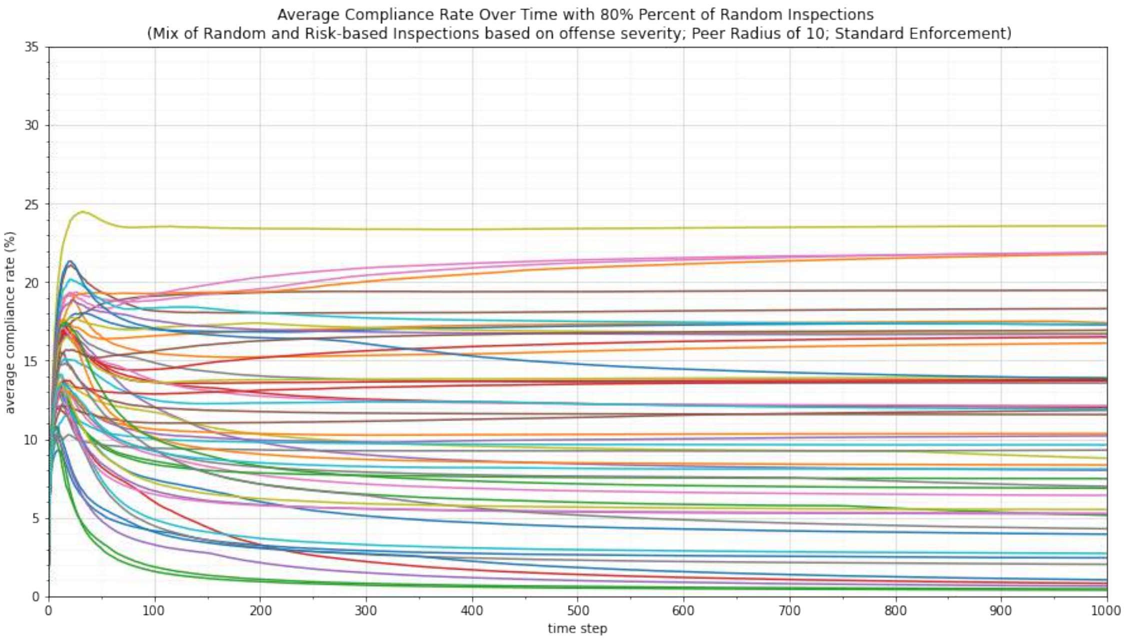

(h) 80% random inspections (SE).

Figure A.26: Average compliance rate results with varying $\%_{rand}$ (SE) (continued).

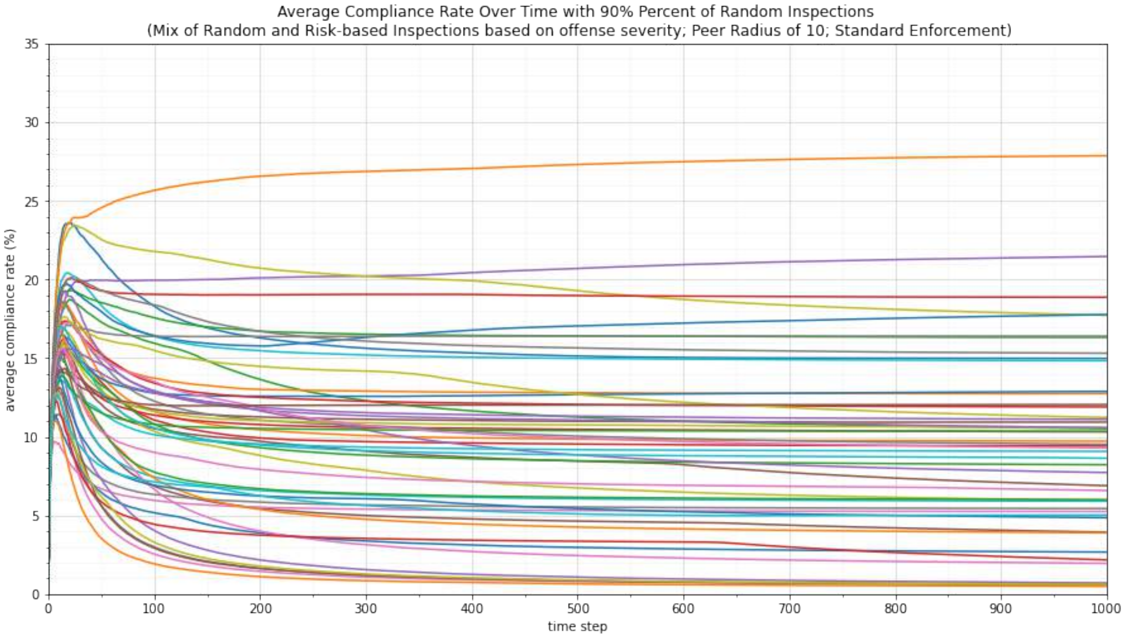

(i) 90% random inspections (SE).

Figure A.26: Average compliance rate results with varying $\%_{rand}$ (SE) (continued).

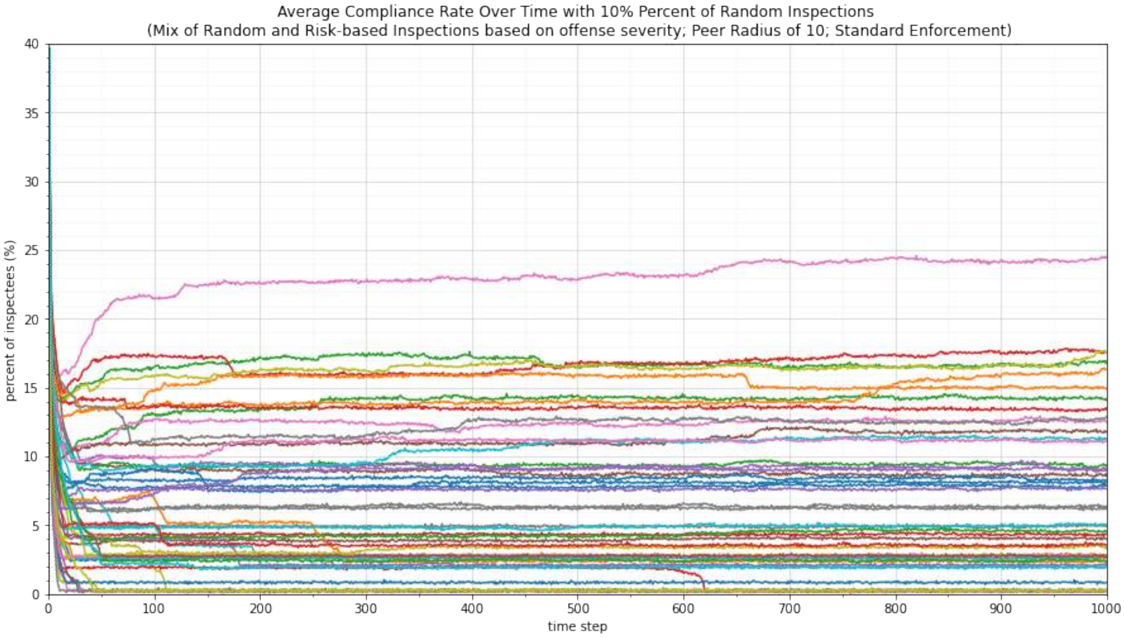

(a) 10% random inspections (SE).

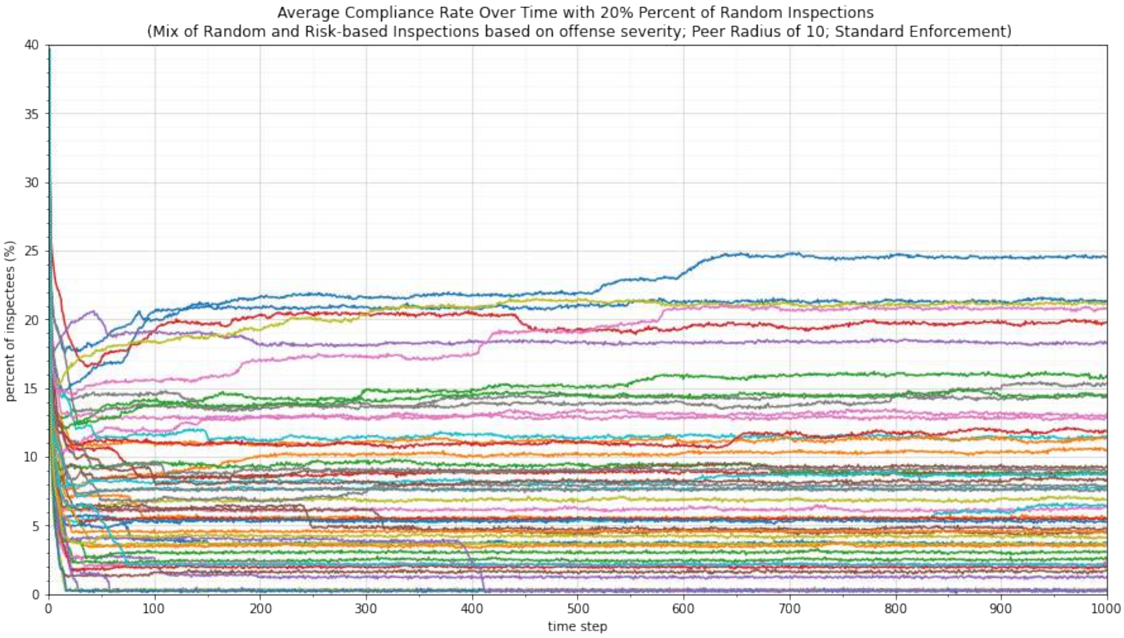

(b) 20% random inspections (SE).

Figure A.27: Share of compliant inspectees with varying $\%_{rand}$ (SE, MRRbOS).

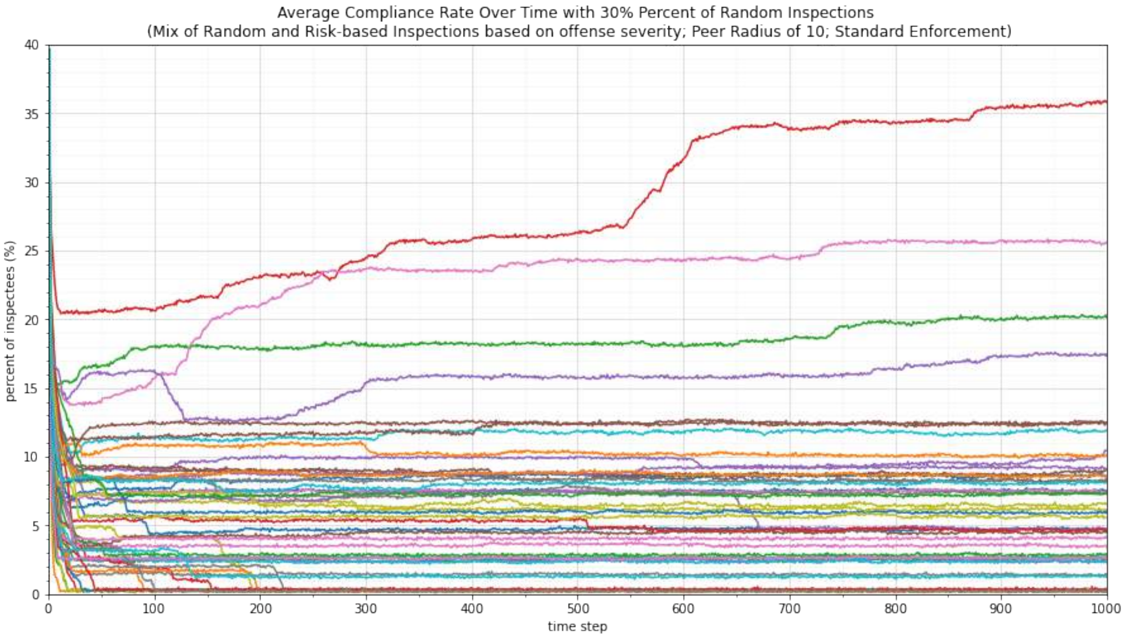

(c) 30% random inspections (SE).

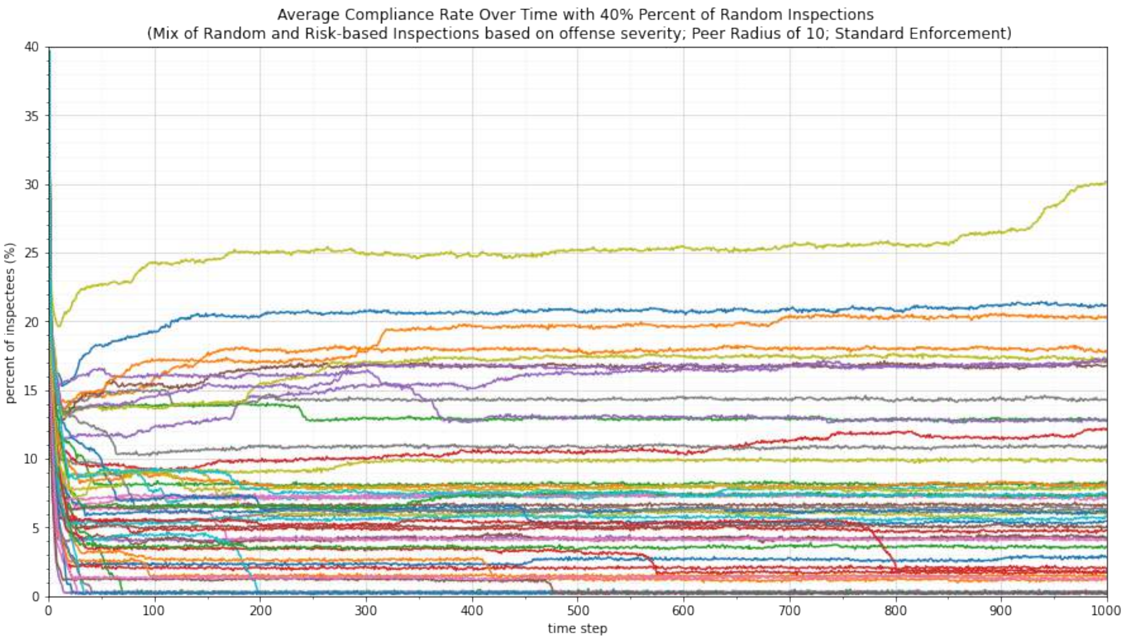

(d) 40% random inspections (SE).

Figure A.27: Share of compliant inspectees with varying $\%_{rand}$ (SE, MRRbOS, continued).

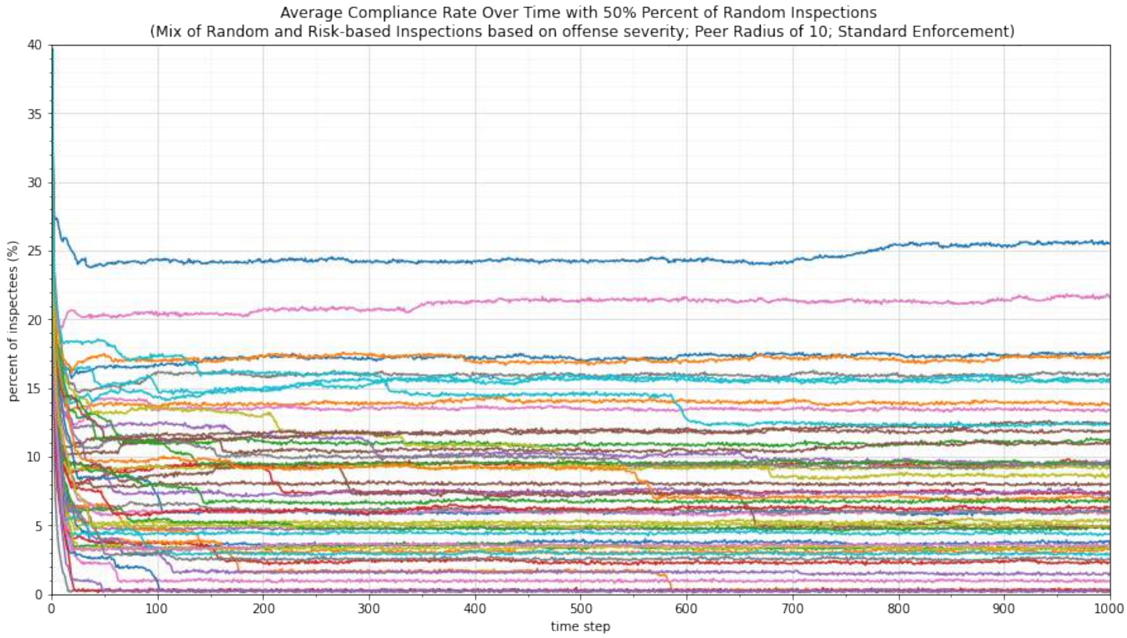

(e) 50% random inspections (SE).

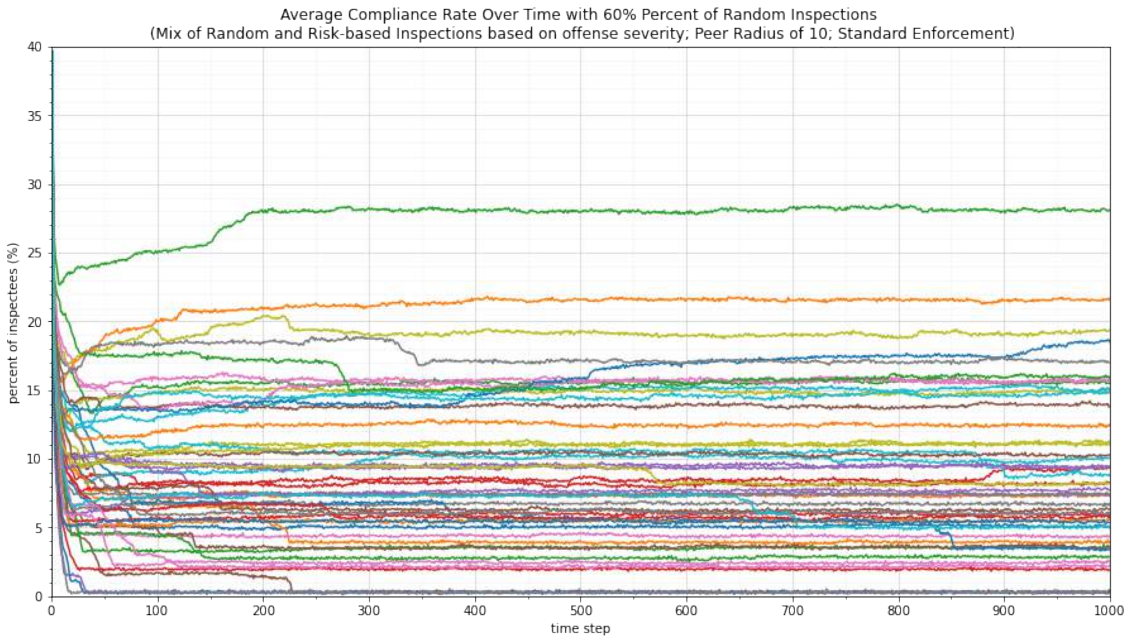

(f) 60% random inspections (SE).

Figure A.27: Share of compliant inspectees with varying $\%_{rand}$ (SE, MRRbOS, continued).

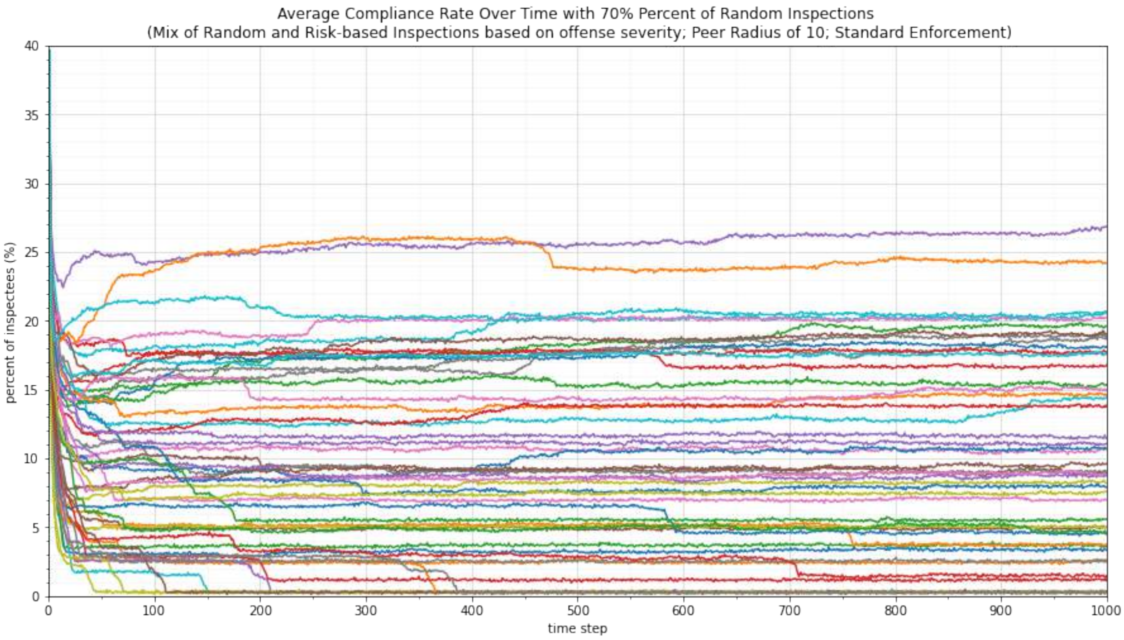

(g) 70% random inspections (SE).

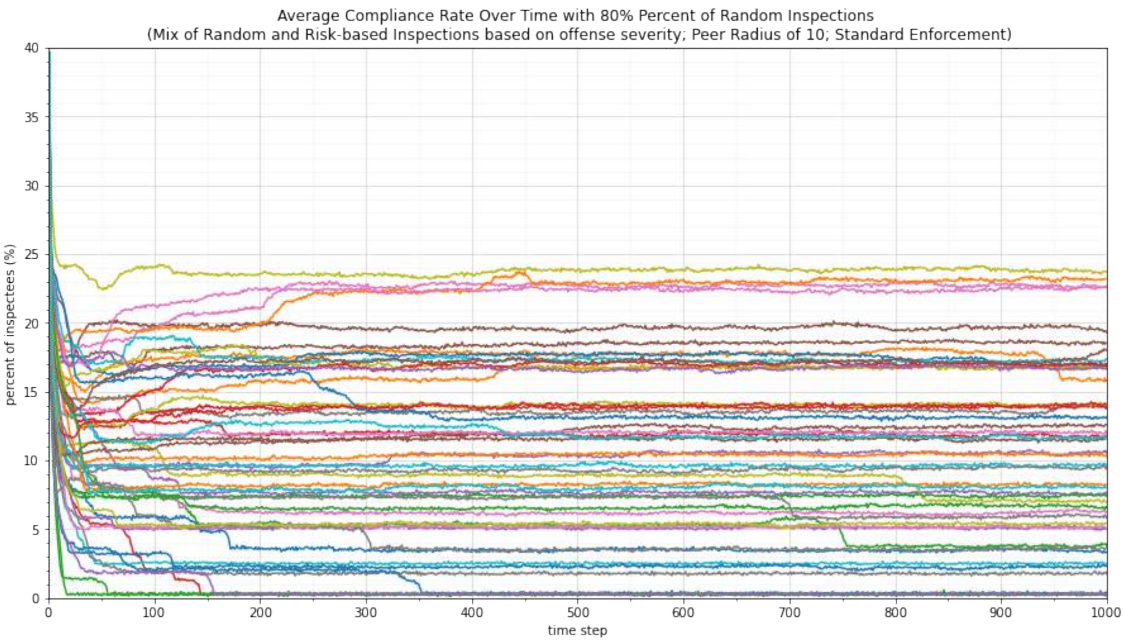

(h) 80% random inspections (SE).

Figure A.27: Share of compliant inspectees with varying $\%_{rand}$ (SE, MRRbOS, continued).

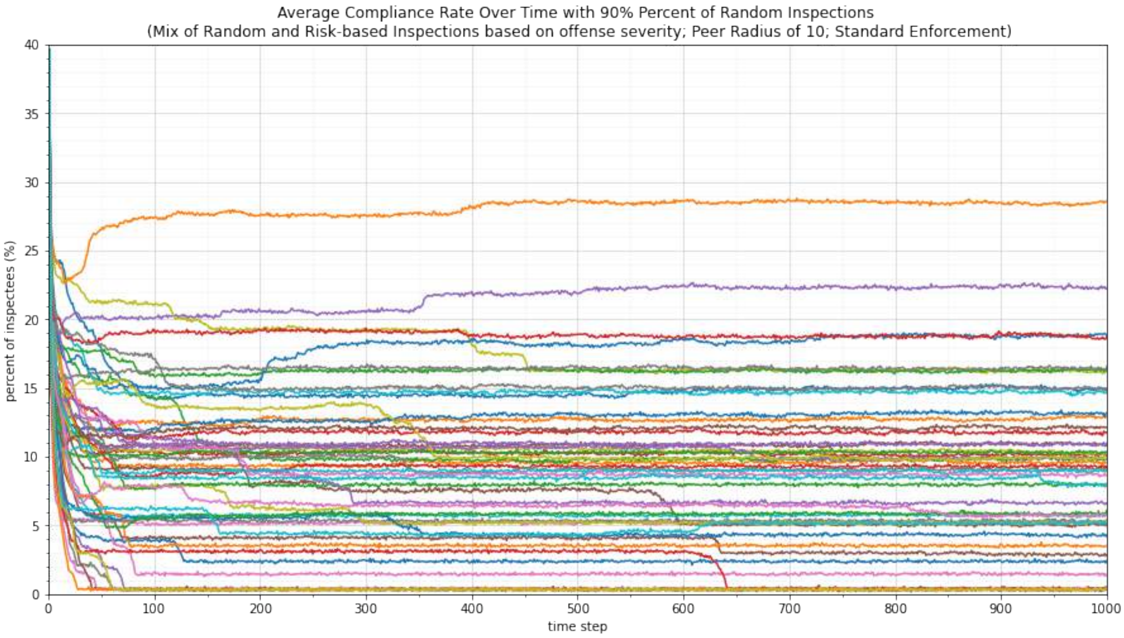

(i) 90% random inspections (SE).

Figure A.27: Share of compliant inspectees with varying $\%_{rand}$ (SE, MRRbOS, continued).

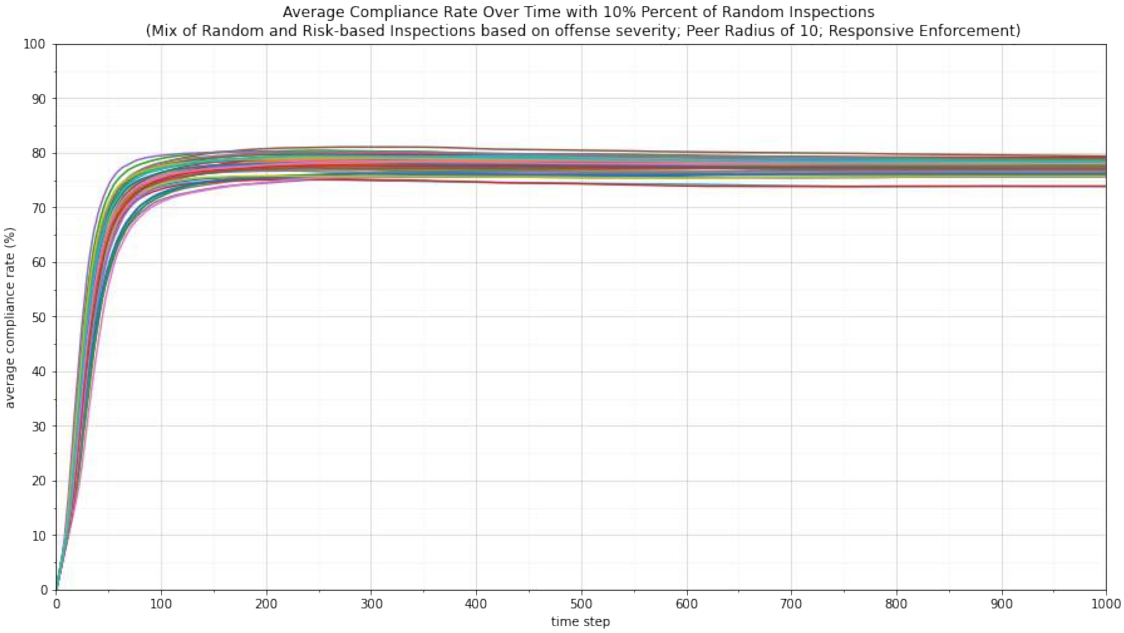

(a) 10% random inspections (RE).

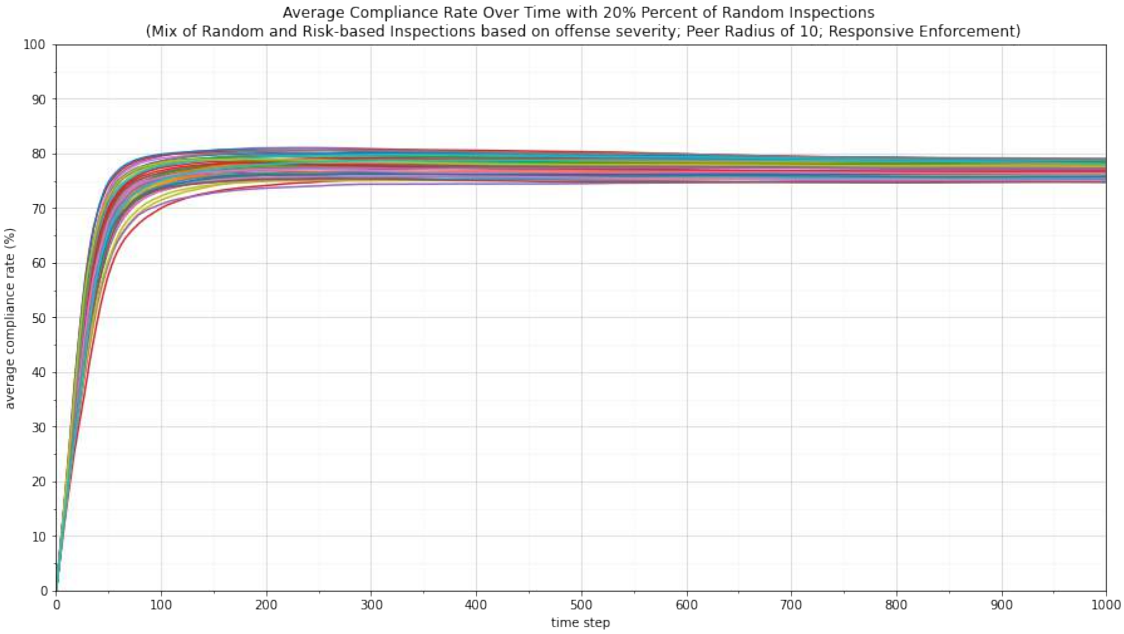

(b) 20% random inspections (RE).

Figure A.28: Average compliance rate results with varying $\%_{rand}$ (RE).

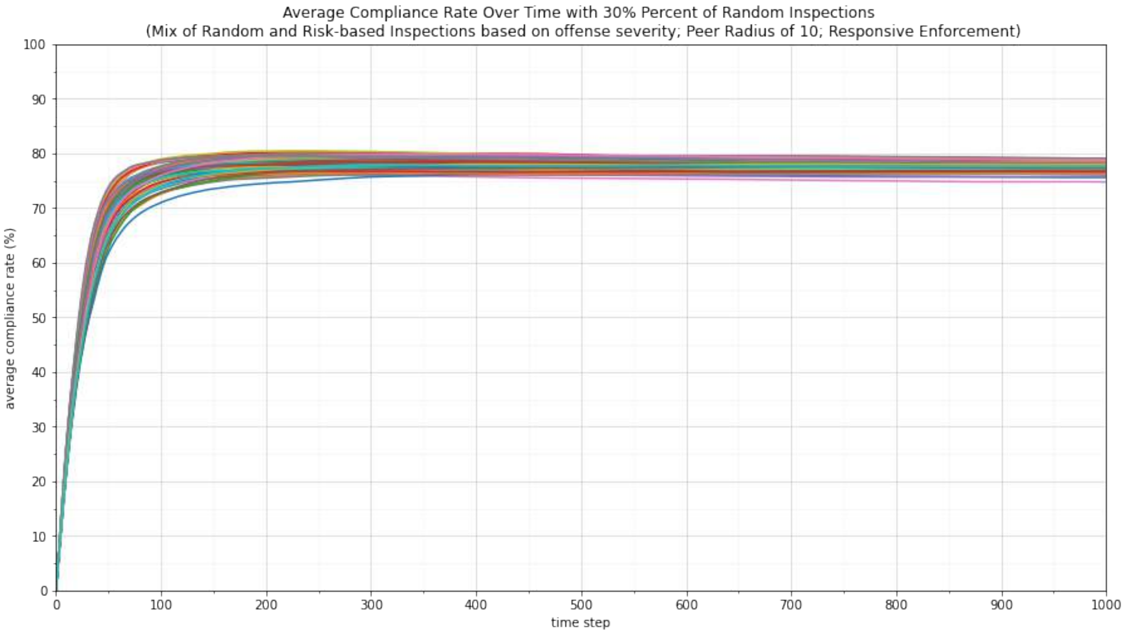

(c) 30% random inspections (RE).

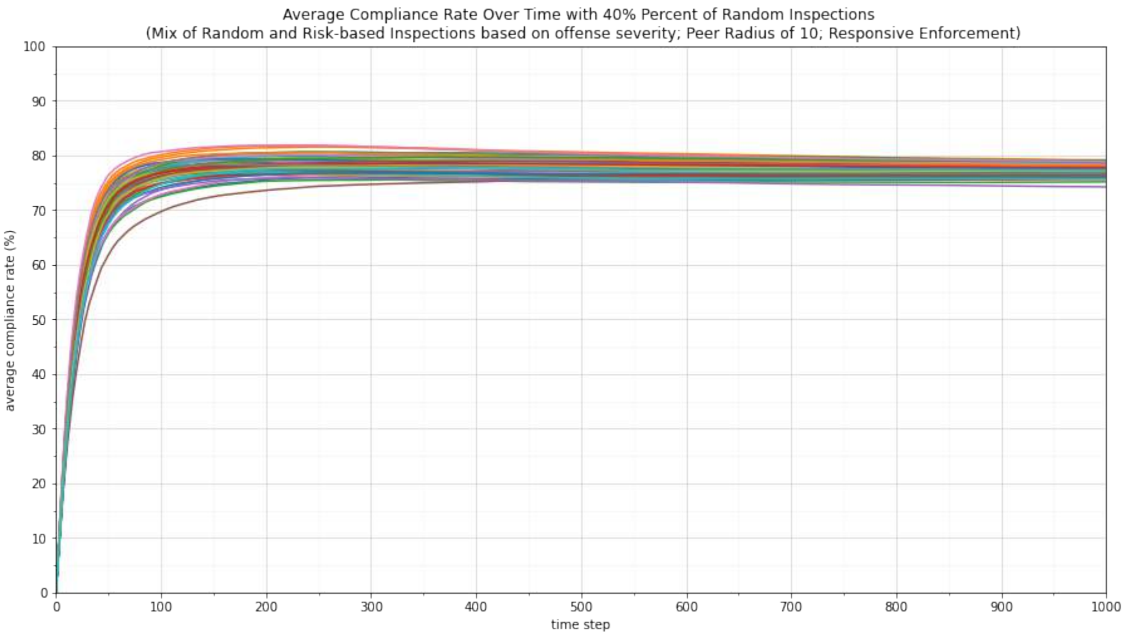

(d) 40% random inspections (RE).

Figure A.28: Average compliance rate results with varying $\%_{rand}$ (RE) (continued).

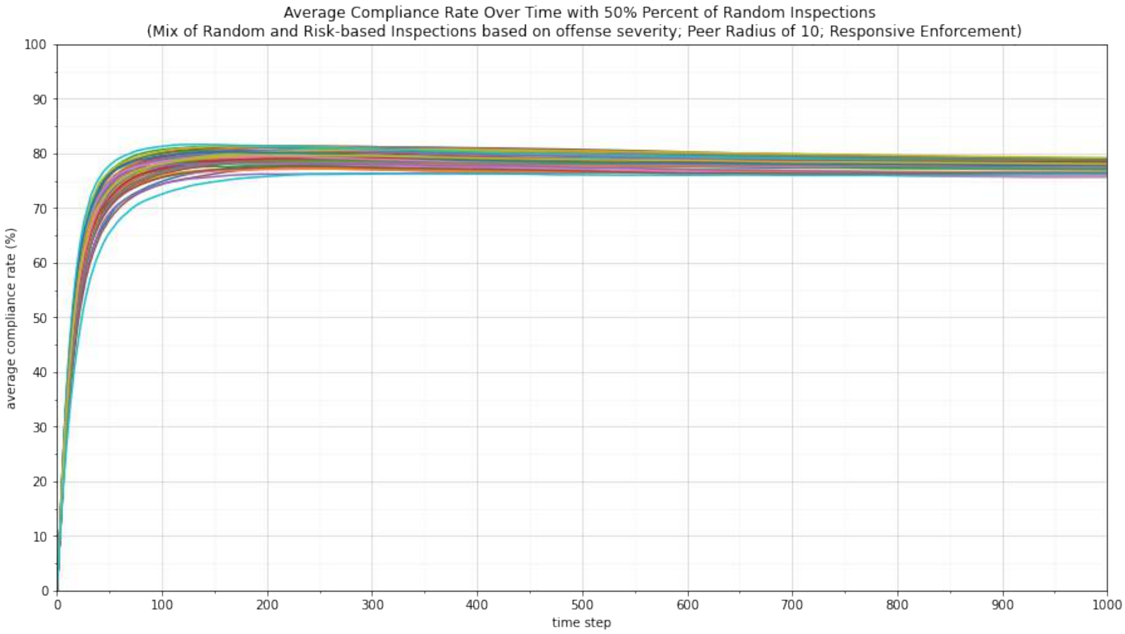

(e) 50% random inspections (RE).

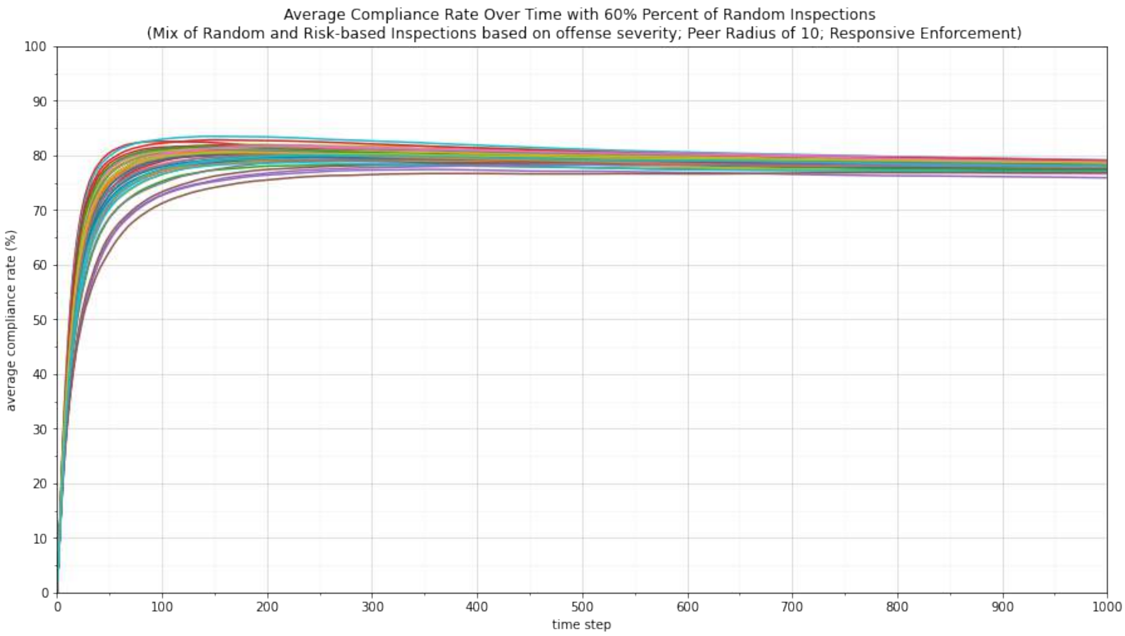

(f) 60% random inspections (RE).

Figure A.28: Average compliance rate results with varying $\%_{rand}$ (RE) (continued).

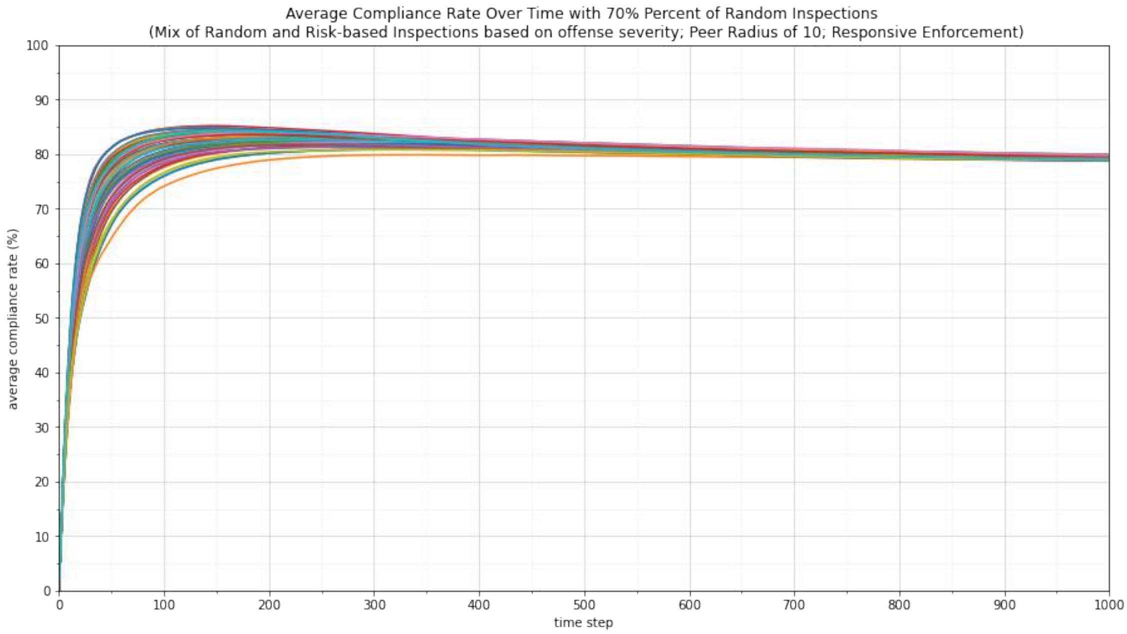

(g) 70% random inspections (RE).

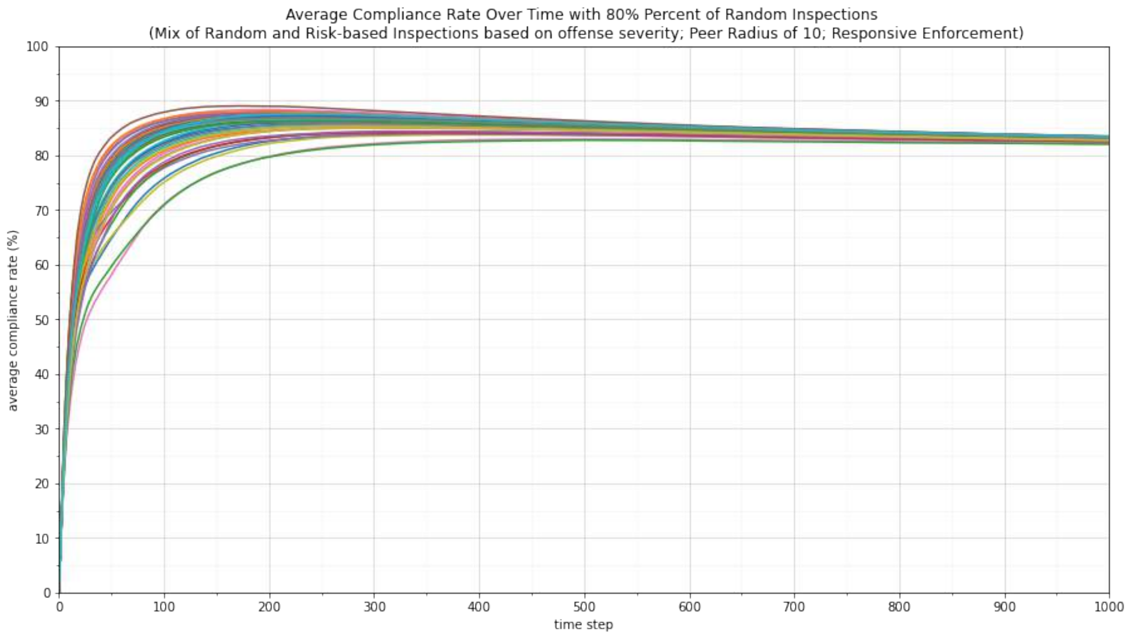

(h) 80% random inspections (RE).

Figure A.28: Average compliance rate results with varying $\%_{rand}$ (RE) (continued).

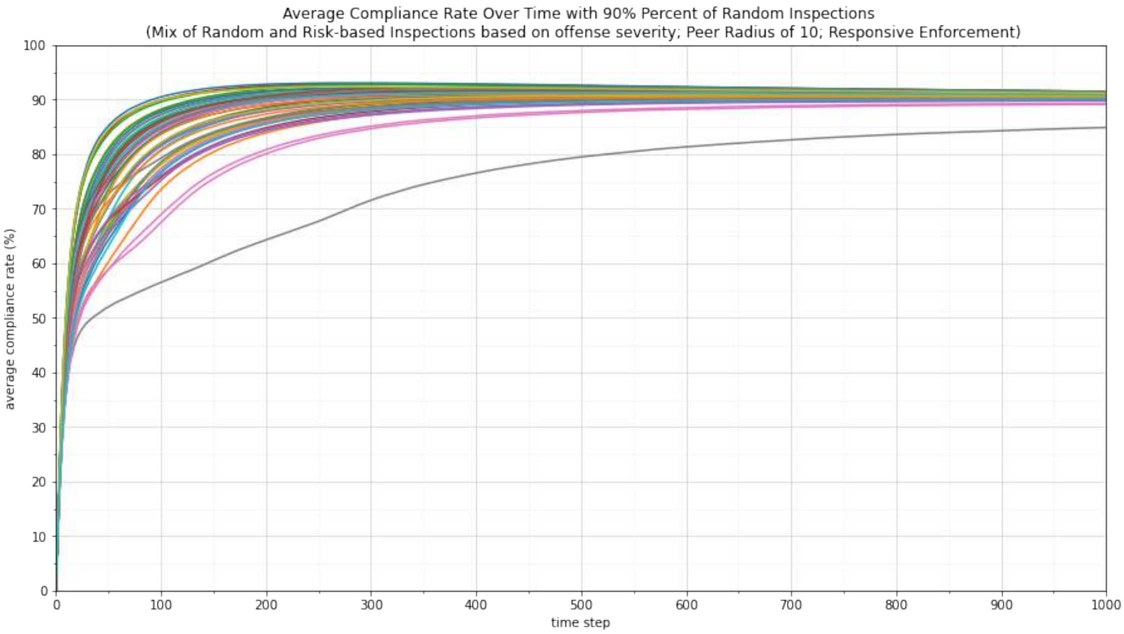

(i) 90% random inspections (RE).

Figure A.28: Average compliance rate results with varying $\%_{rand}$ (RE) (continued).

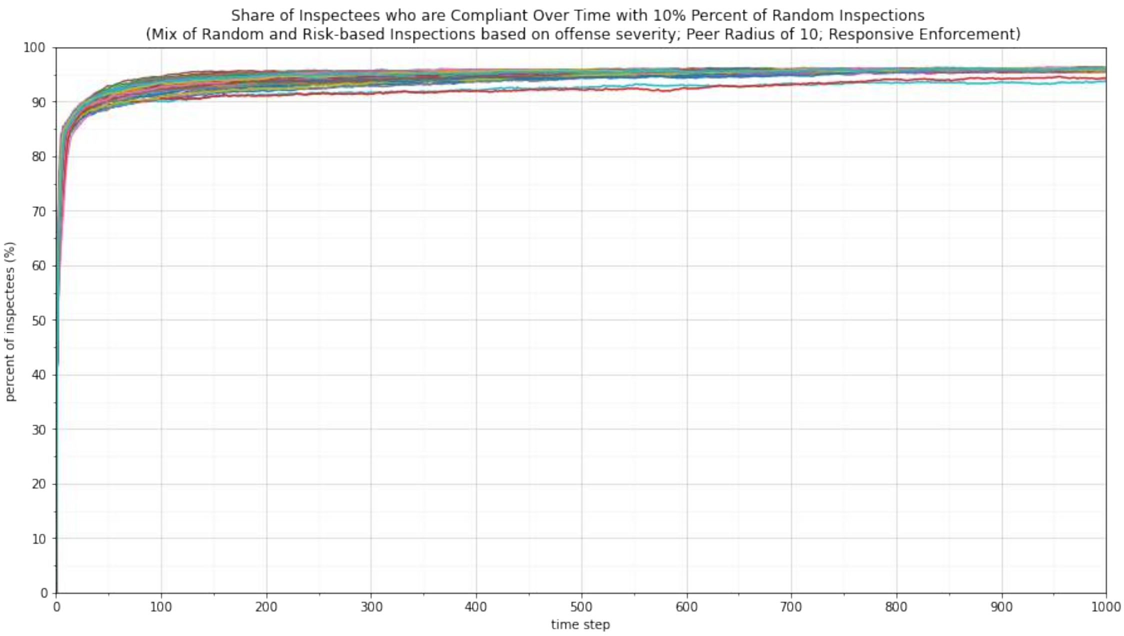

(a) 10% random inspections (RE).

Figure A.29: Share of compliant inspectees with varying $\%_{rand}$ (RE, MRRbOS).

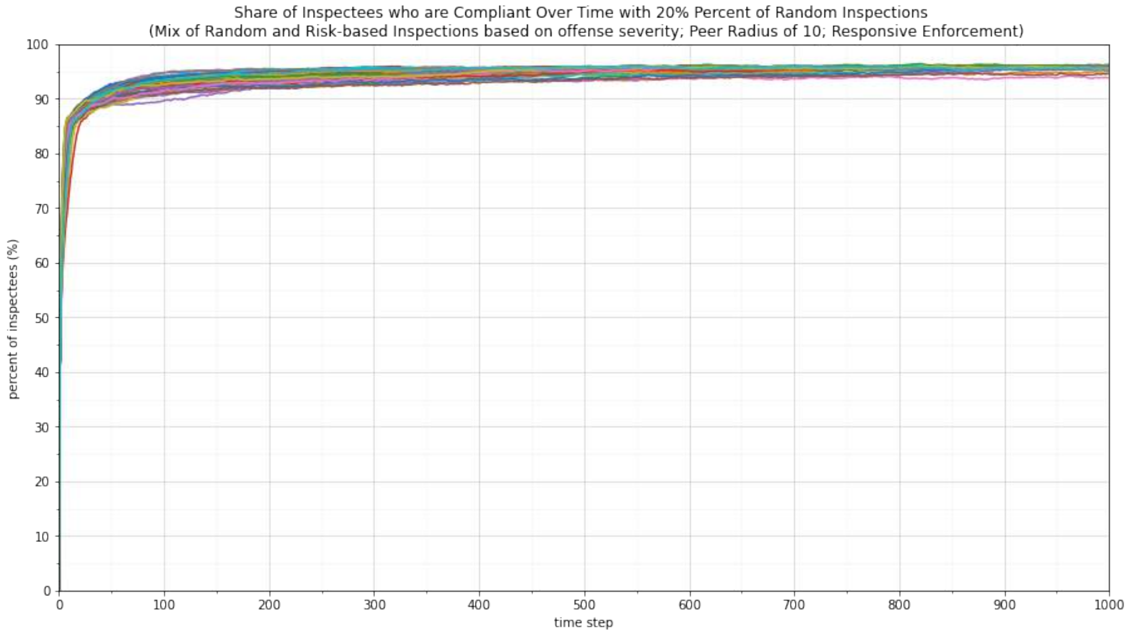

(b) 20% random inspections (RE).

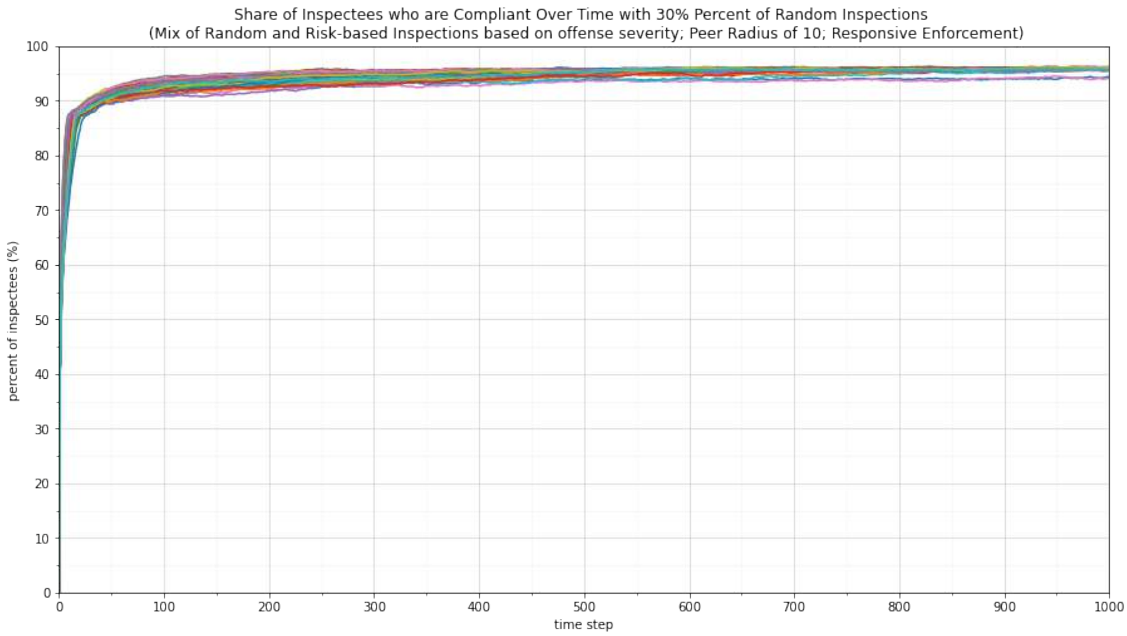

(c) 30% random inspections (RE).

Figure A.29: Share of compliant inspectees with varying $\%_{rand}$ (RE, MRRbOS, continued).

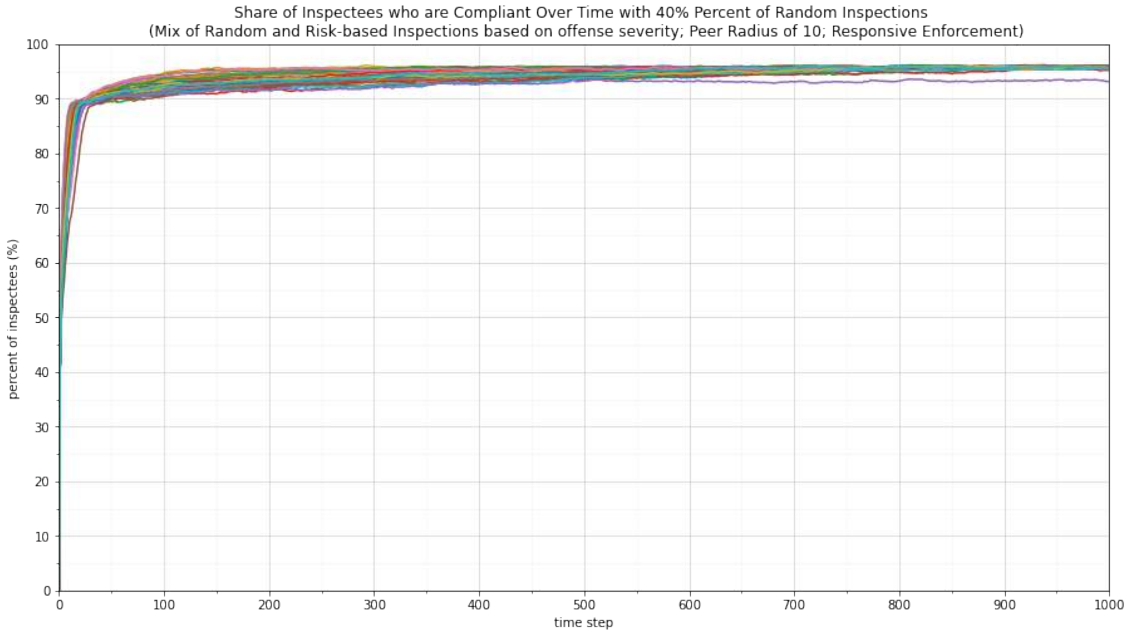

(d) 40% random inspections (RE).

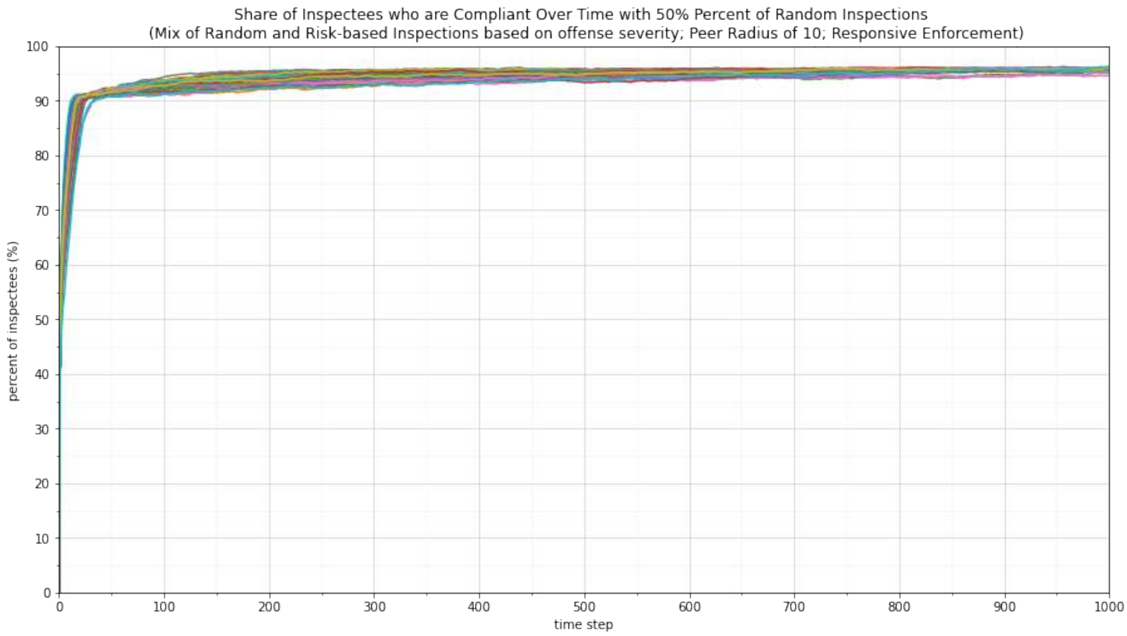

(e) 50% random inspections (RE).

Figure A.29: Share of compliant inspectees with varying $\%_{rand}$ (RE, MRRbOS, continued).

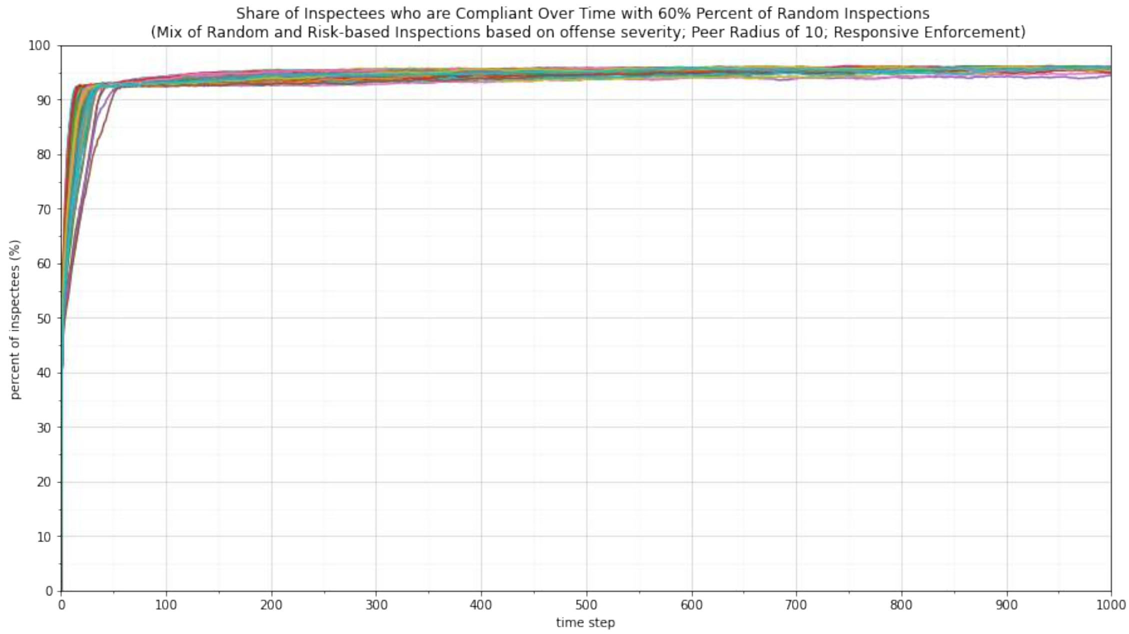

(f) 60% random inspections (RE).

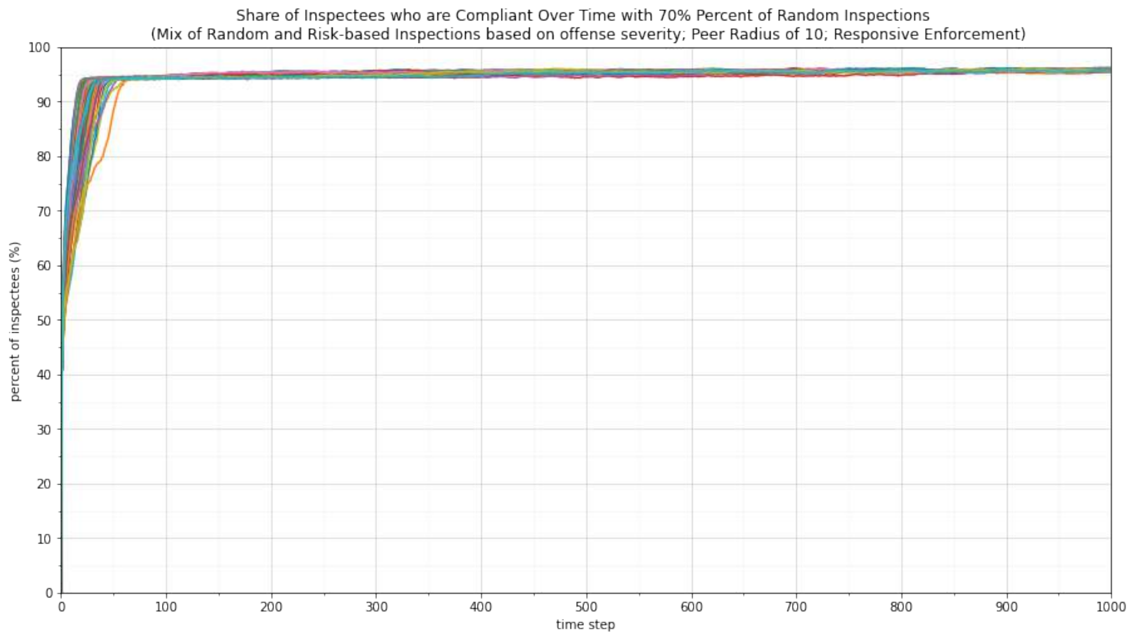

(g) 70% random inspections (RE).

Figure A.29: Share of compliant inspectees with varying $\%_{rand}$ (RE, MRRbOS, continued).

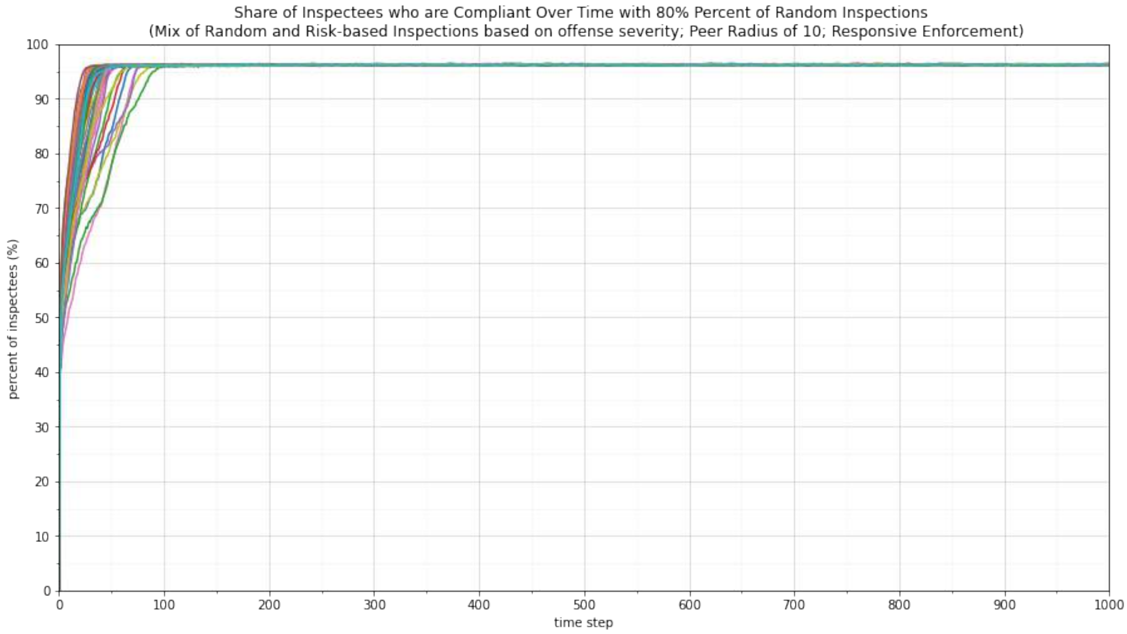

(h) 80% random inspections (RE).

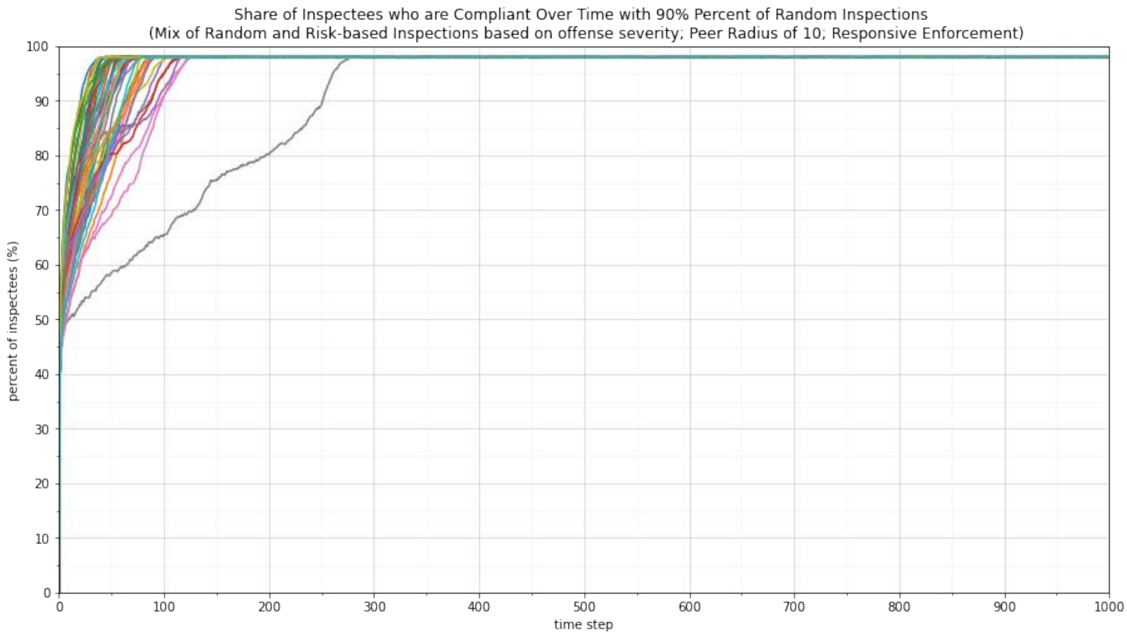

(i) 90% random inspections (RE).

Figure A.29: Share of compliant inspectees with varying $\%_{rand}$ (RE, MRRbOS, continued).

A.1.3 Varying Logit Parameters (μ and σ)

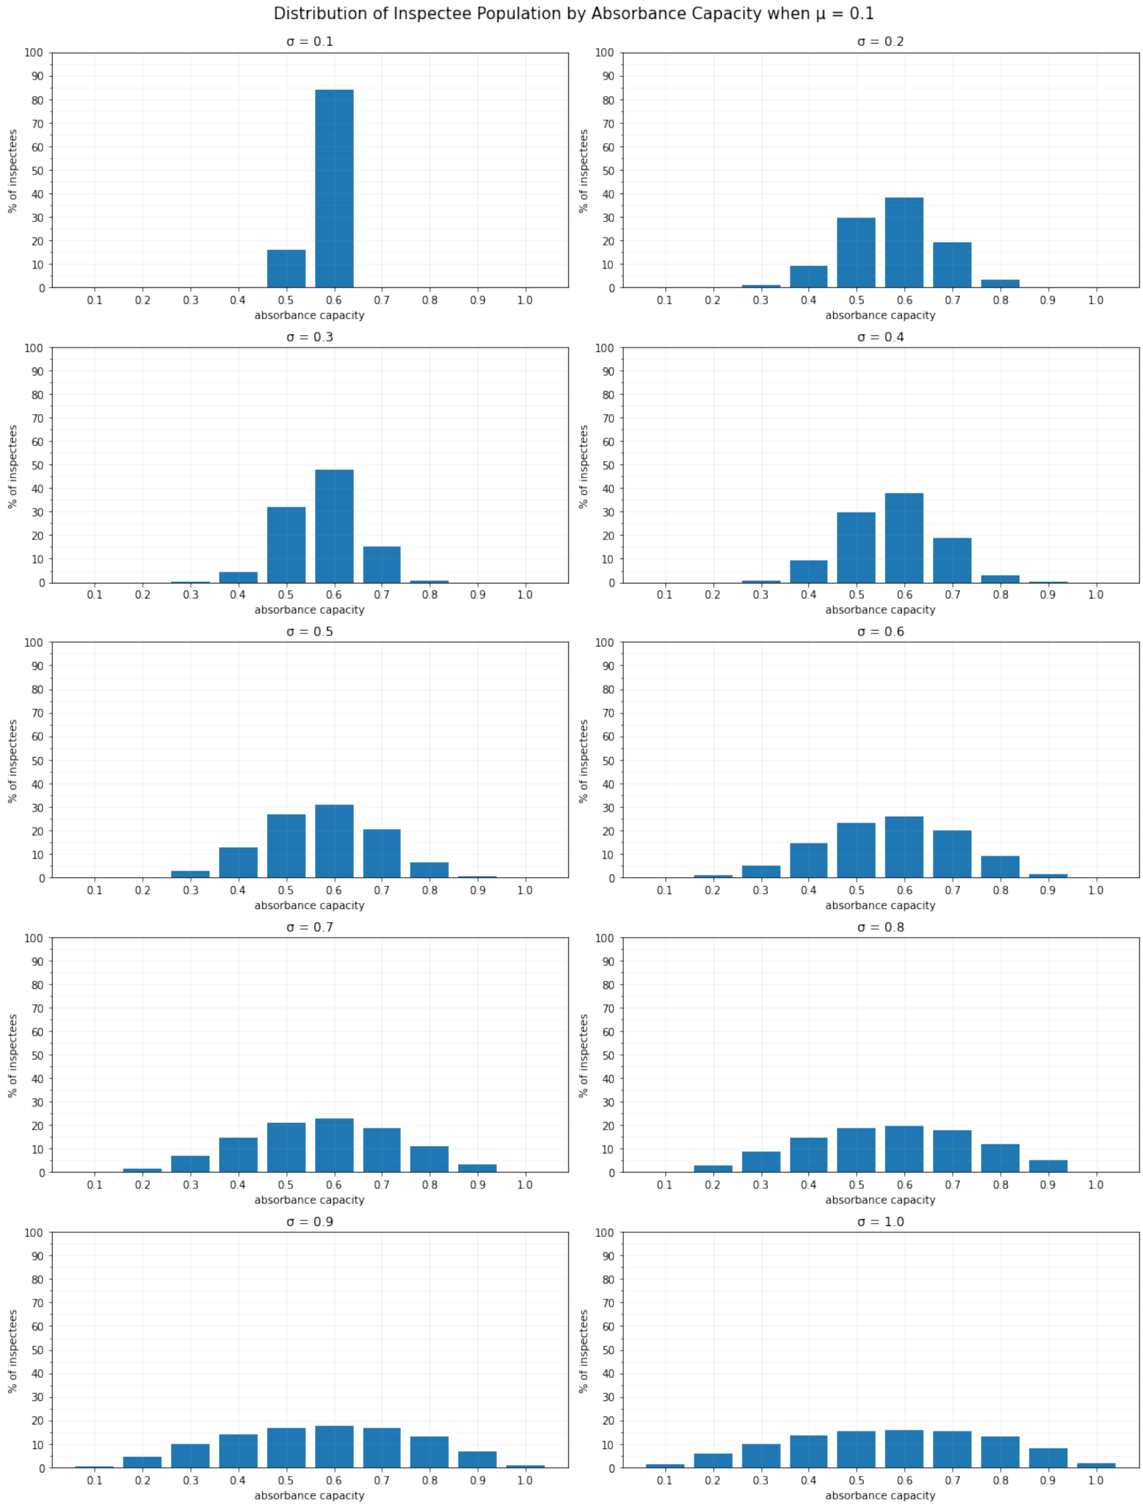

Figure A.30: Inspectee population distribution by absorbance capacity with μ=0.1.

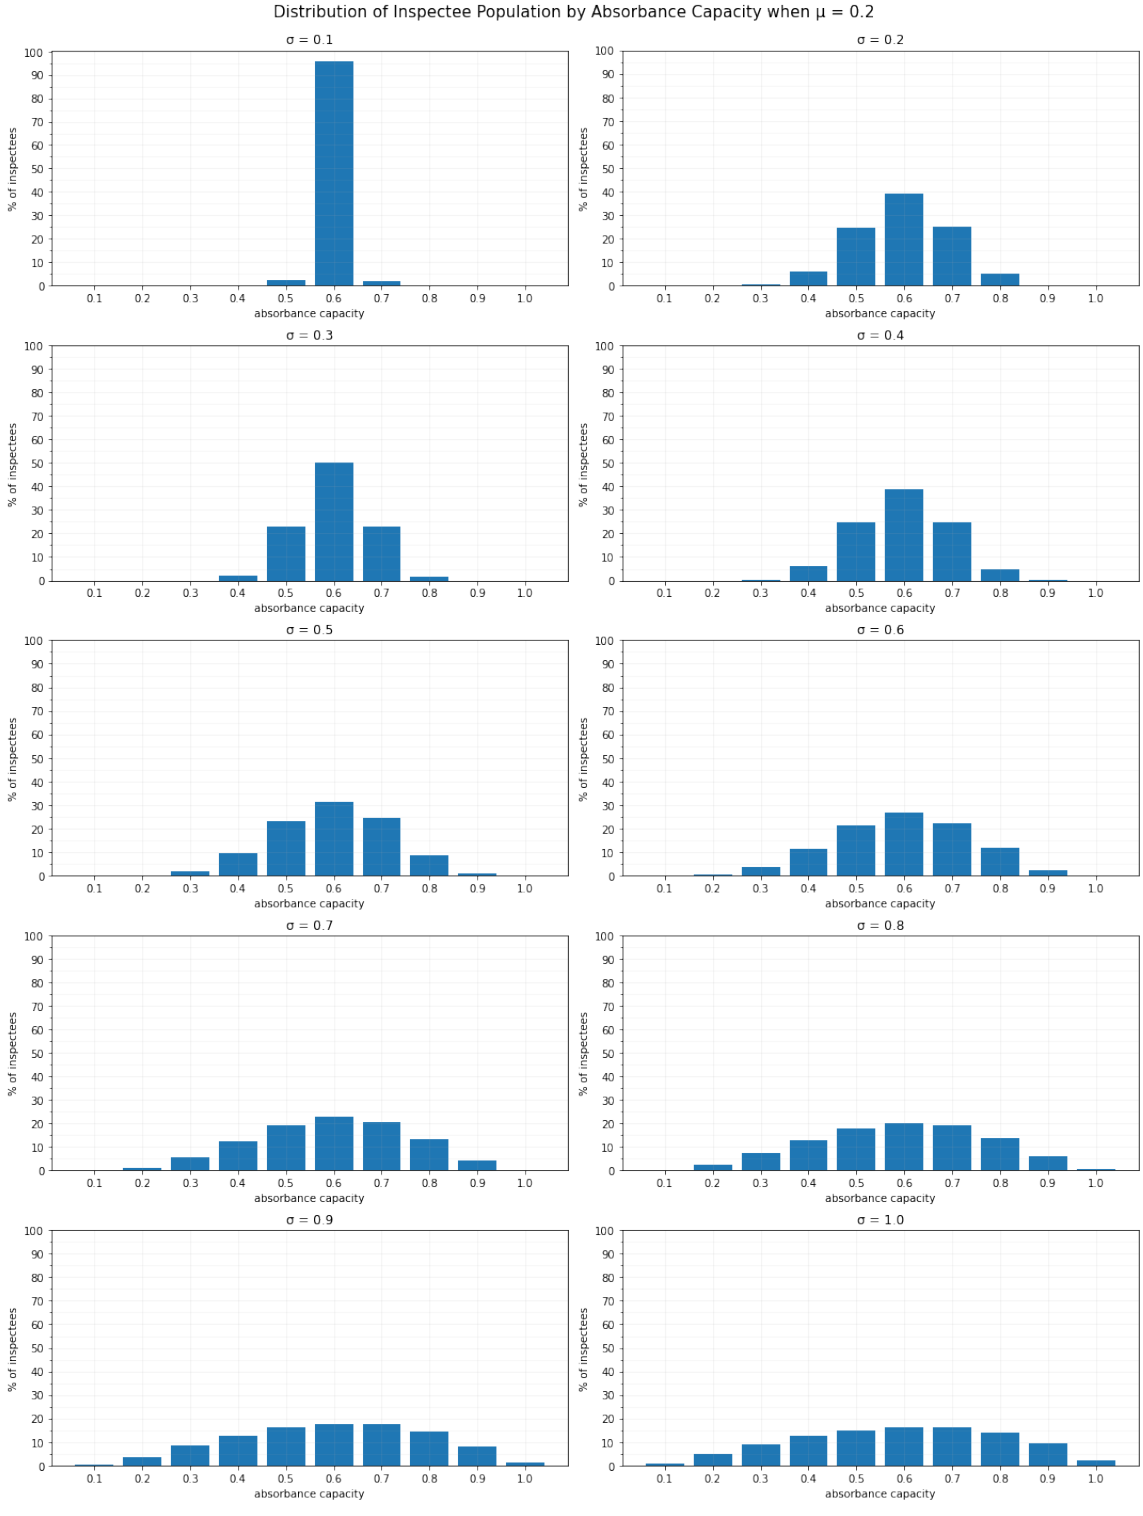

Figure A.31: Inspectee population distribution by absorbance capacity with μ=0.2.

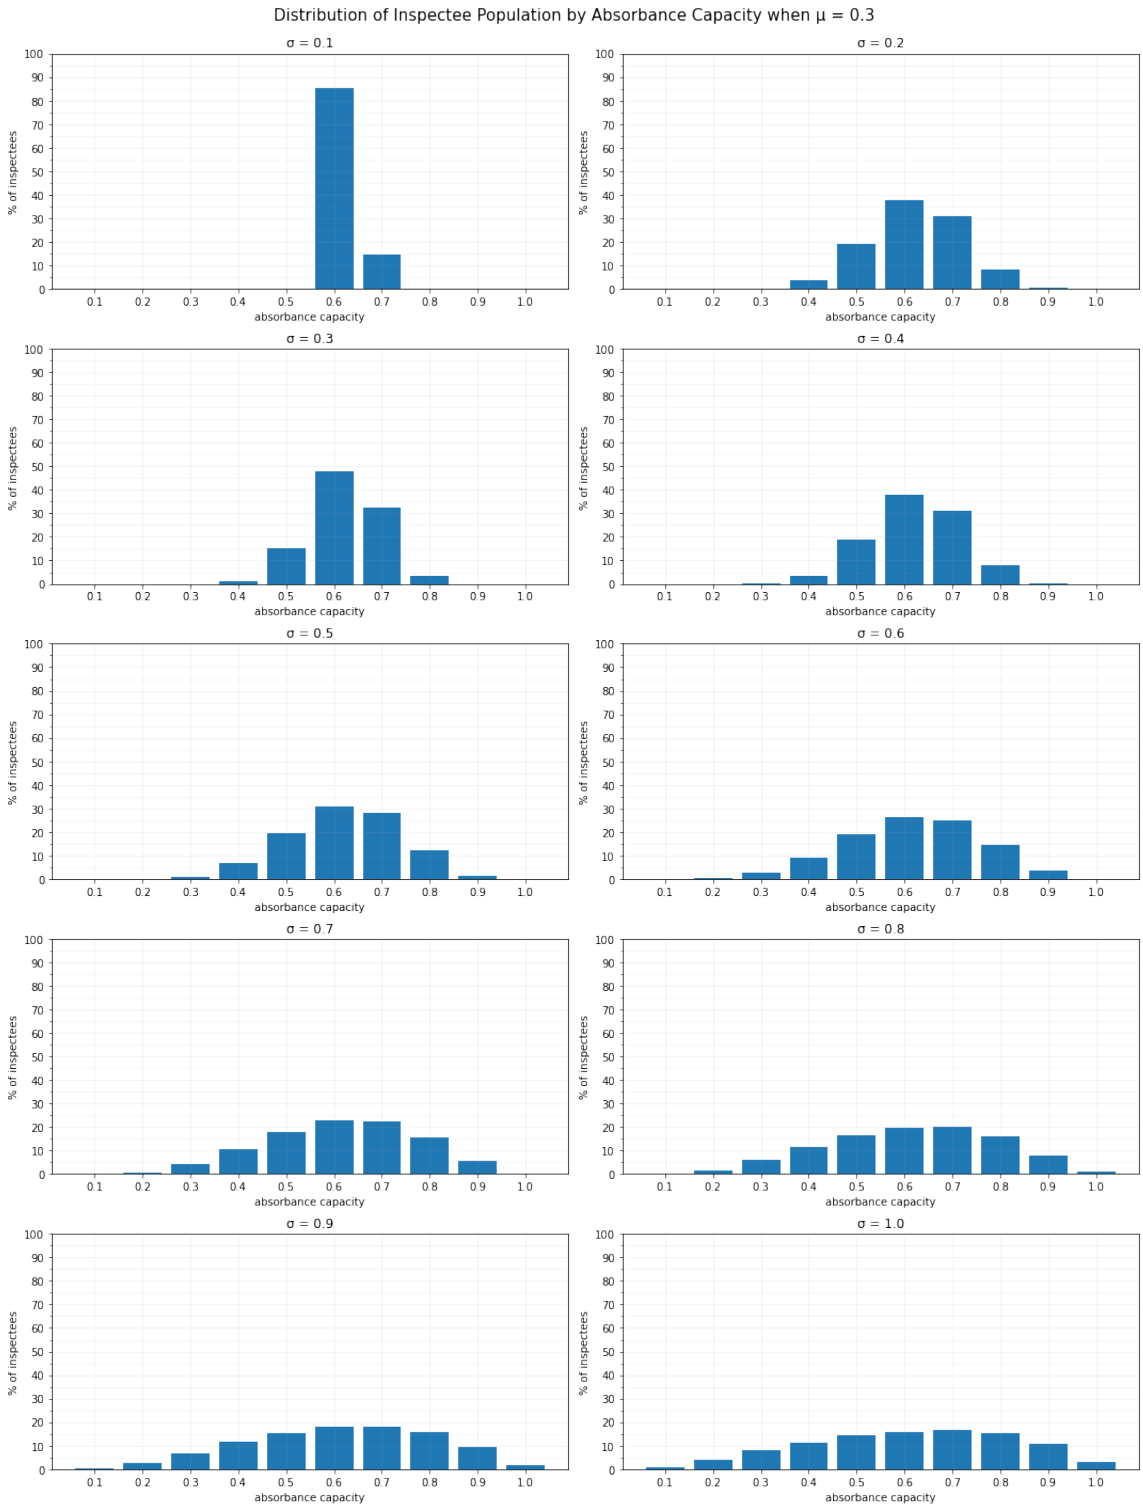

Figure A.32: Inspectee population distribution by absorbance capacity with μ=0.3.

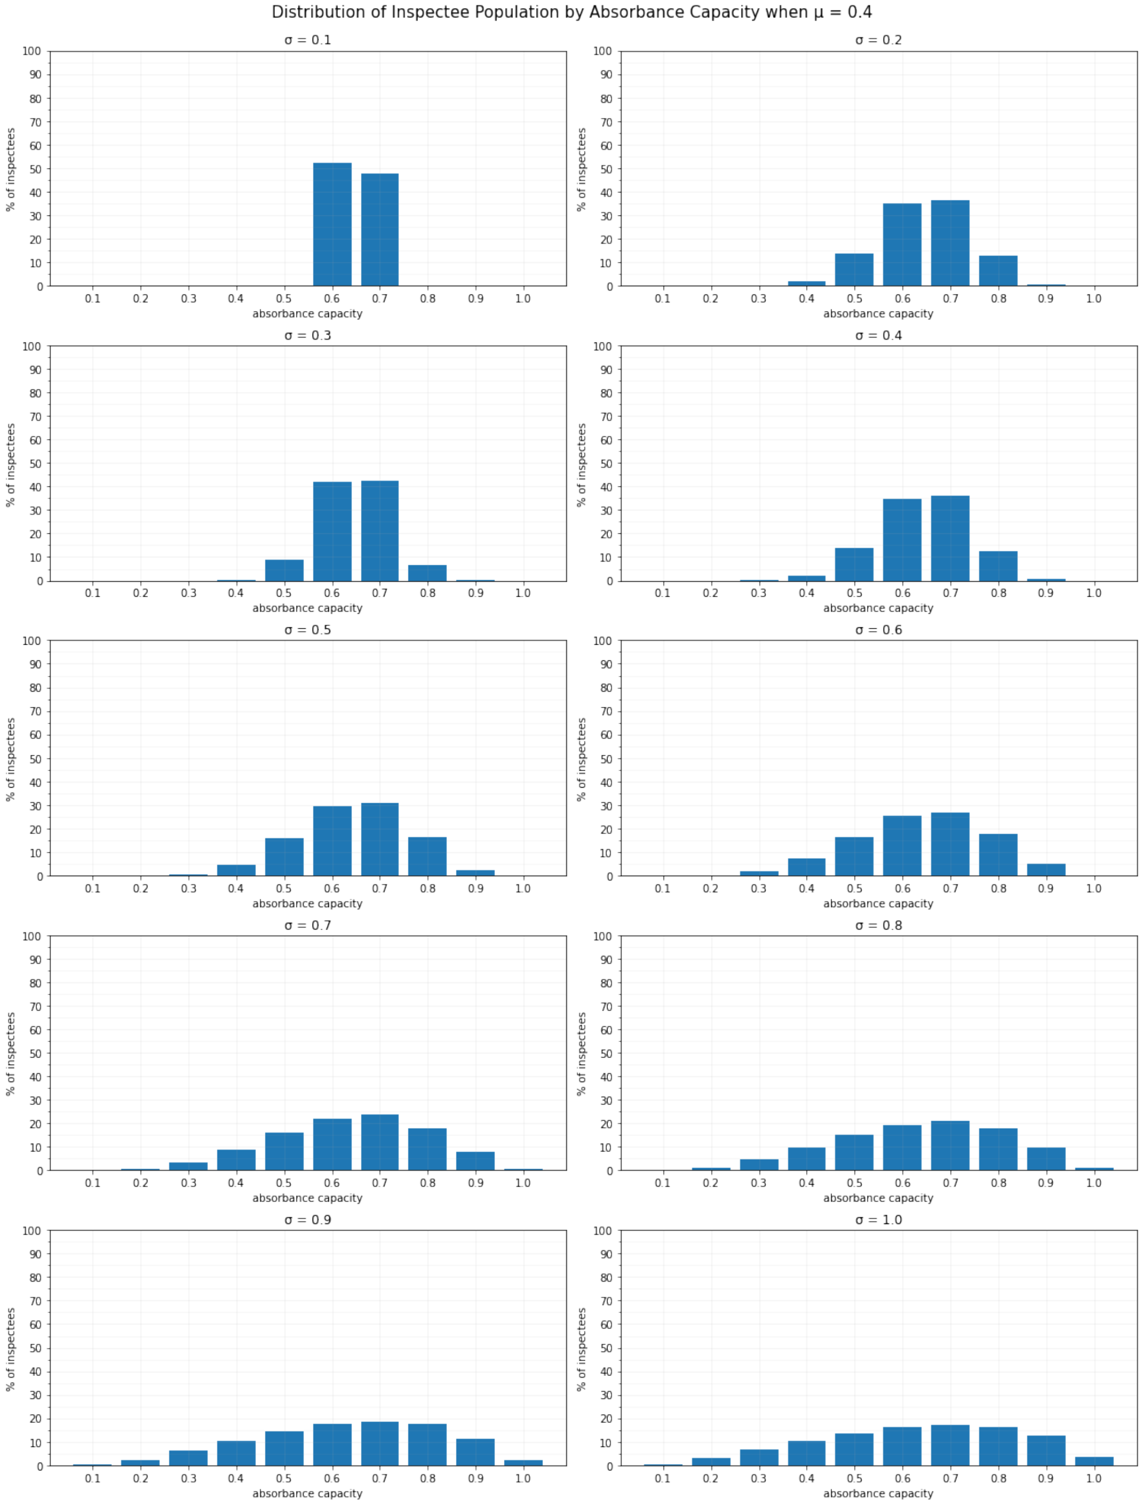

Figure A.33: Inspectee population distribution by absorbance capacity with μ=0.4.

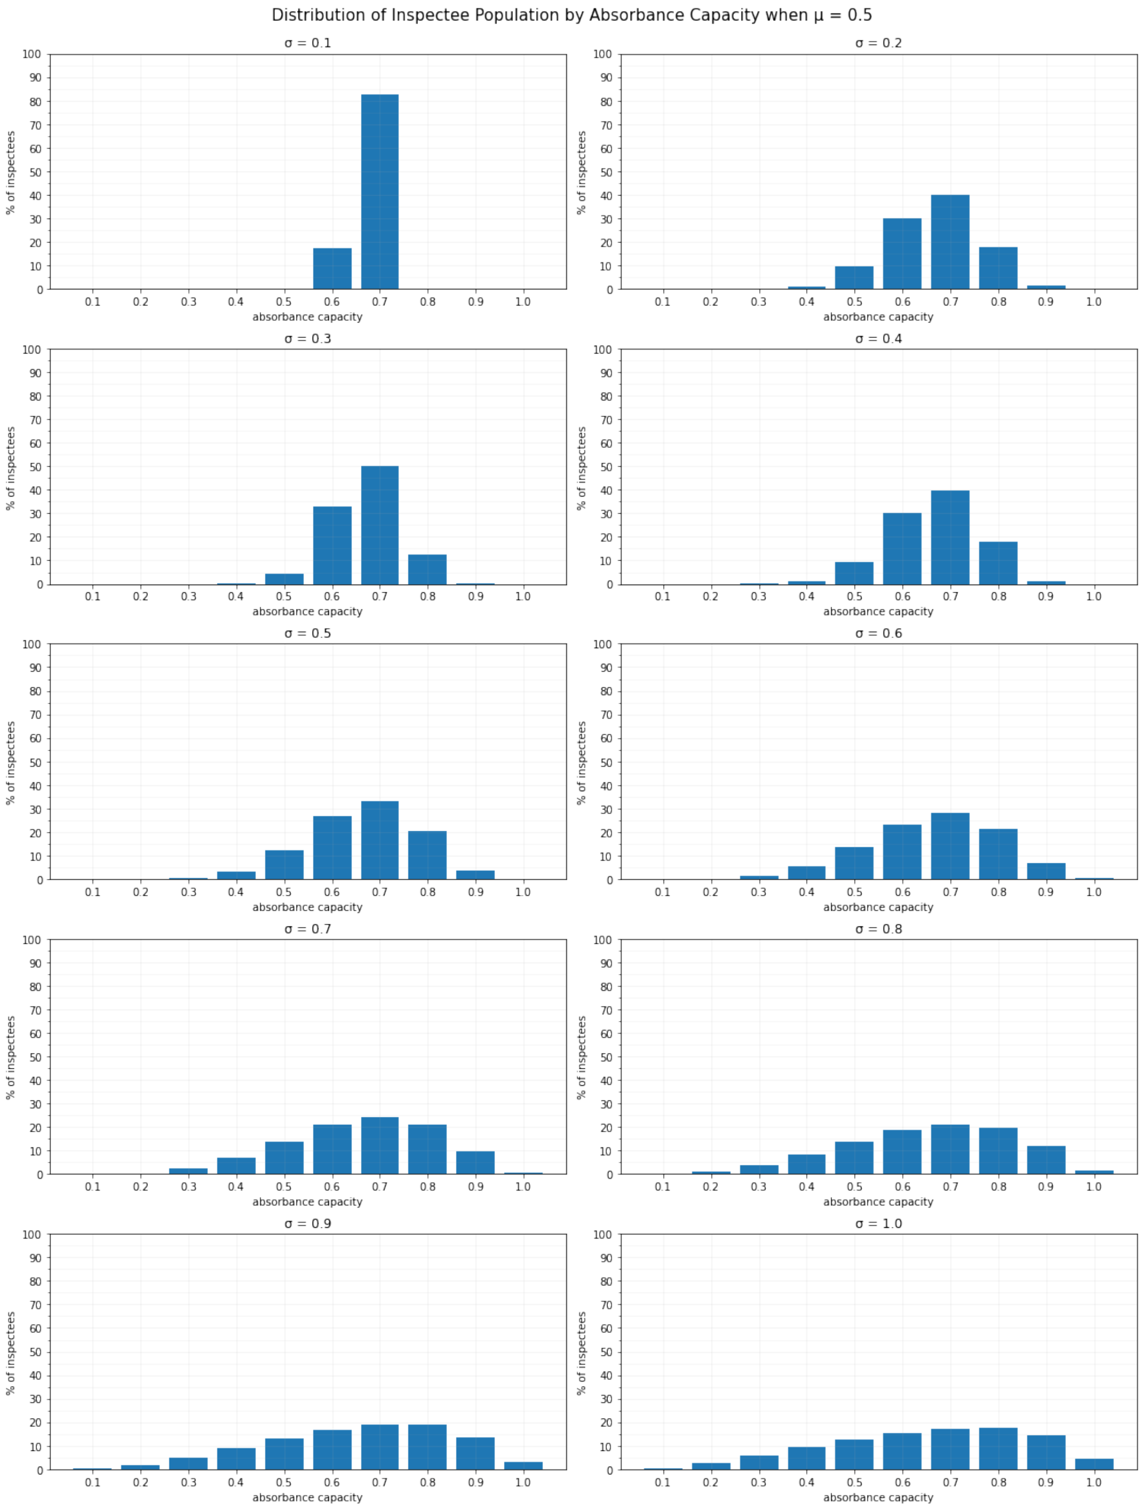

Figure A.34: Inspectee population distribution by absorbance capacity with μ=0.5.

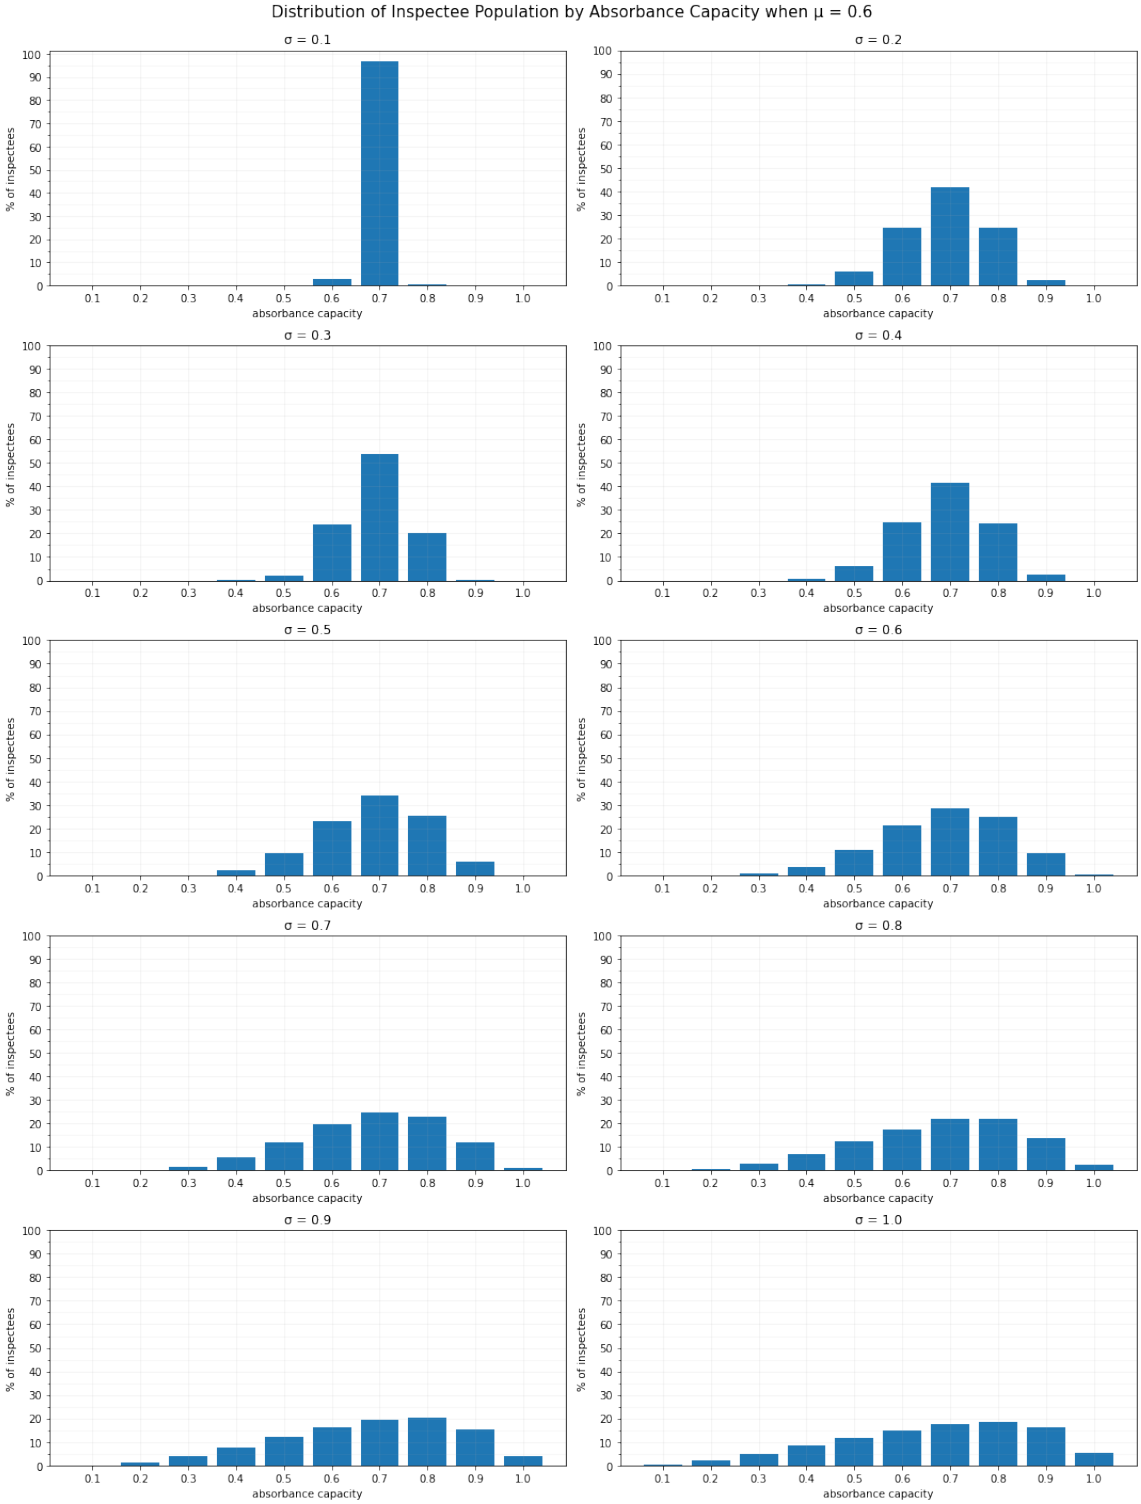

Figure A.35: Inspectee population distribution by absorbance capacity with μ=0.6.

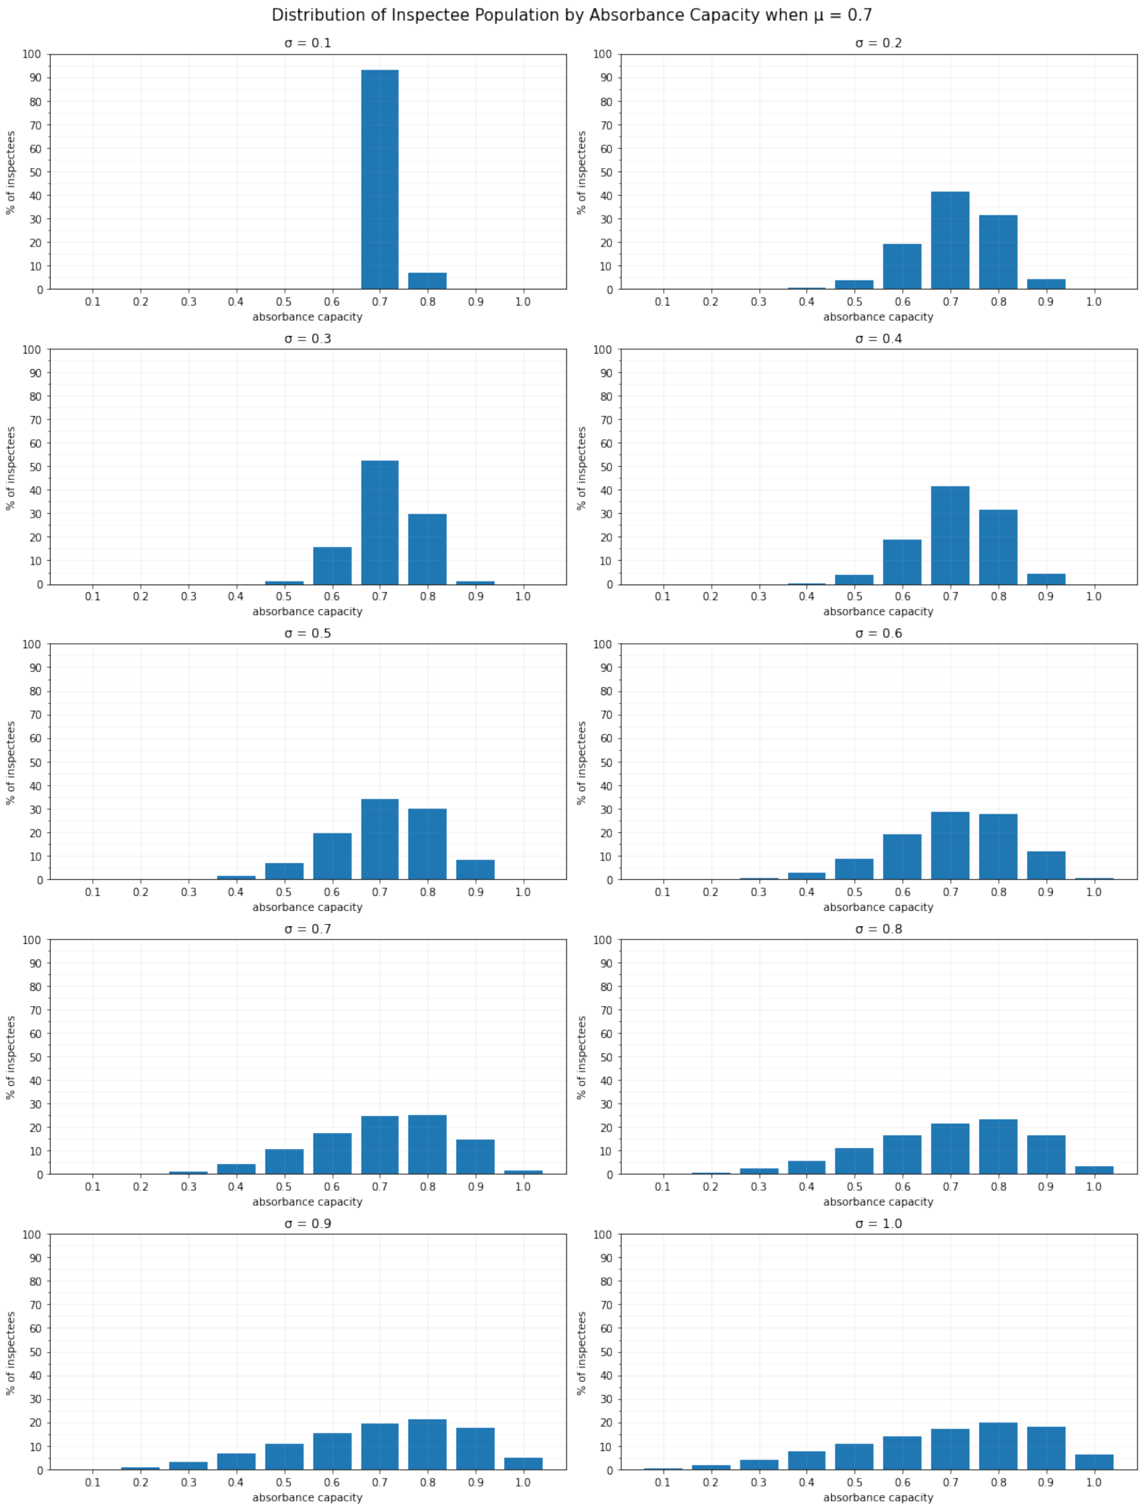

Figure A.36: Inspectee population distribution by absorbance capacity with μ=0.7.

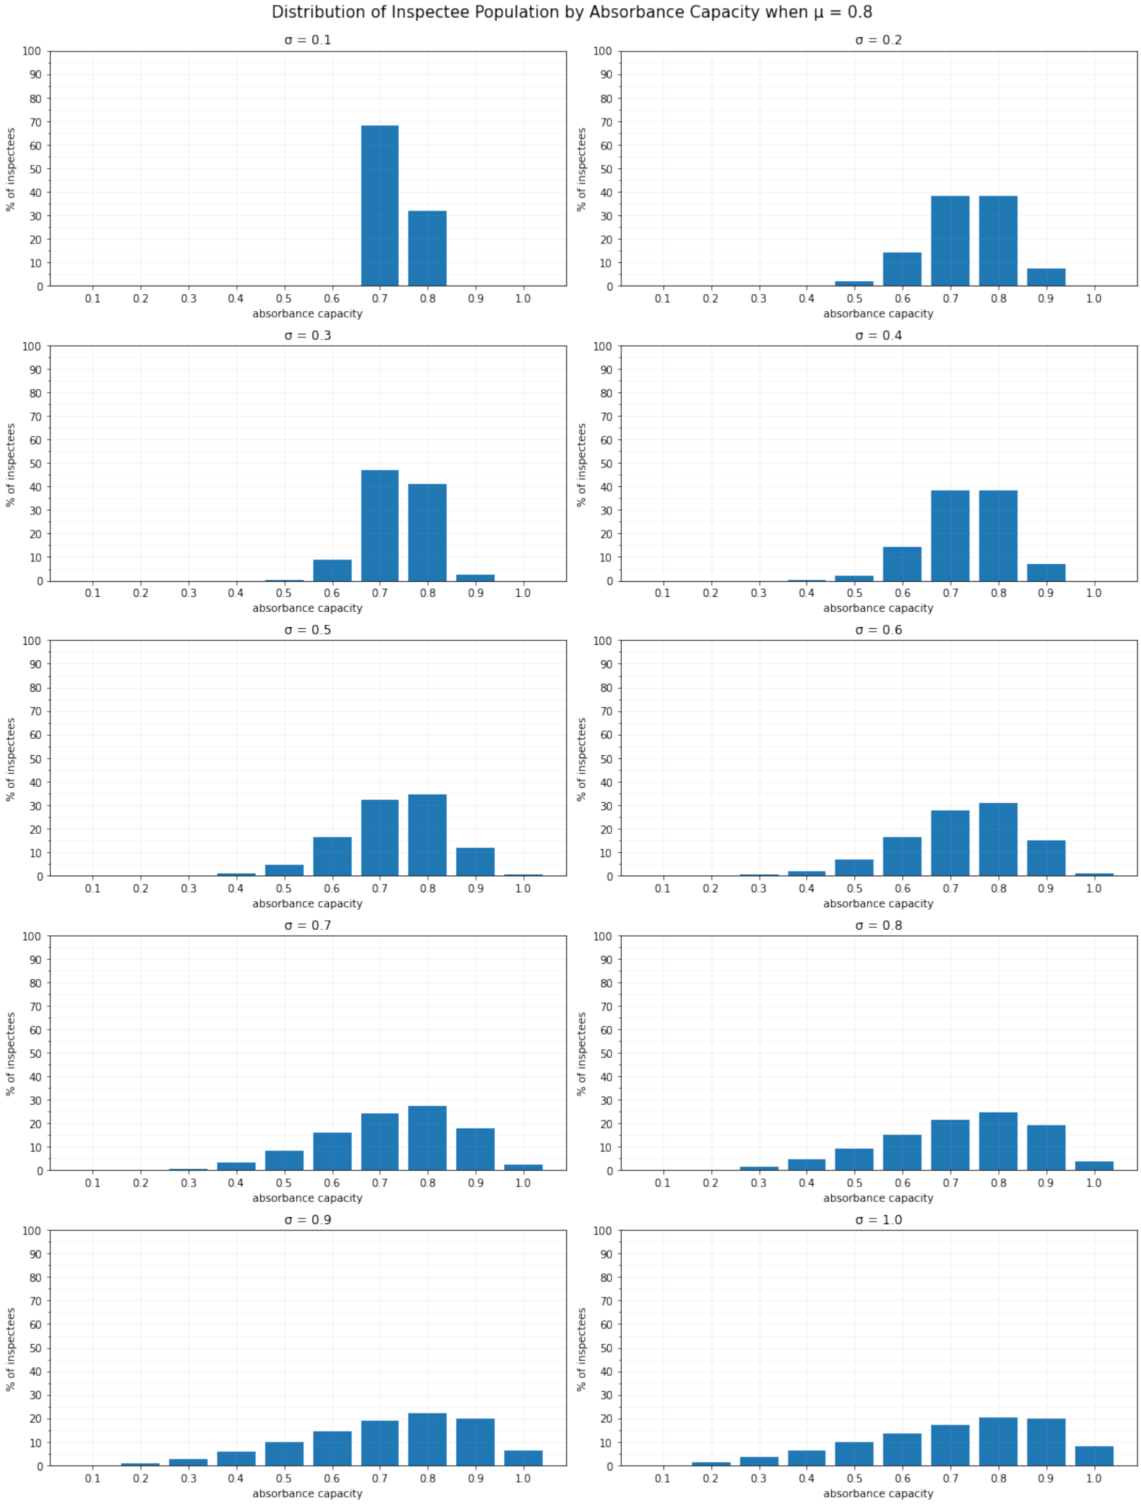

Figure A.37: Inspectee population distribution by absorbance capacity with μ=0.8.

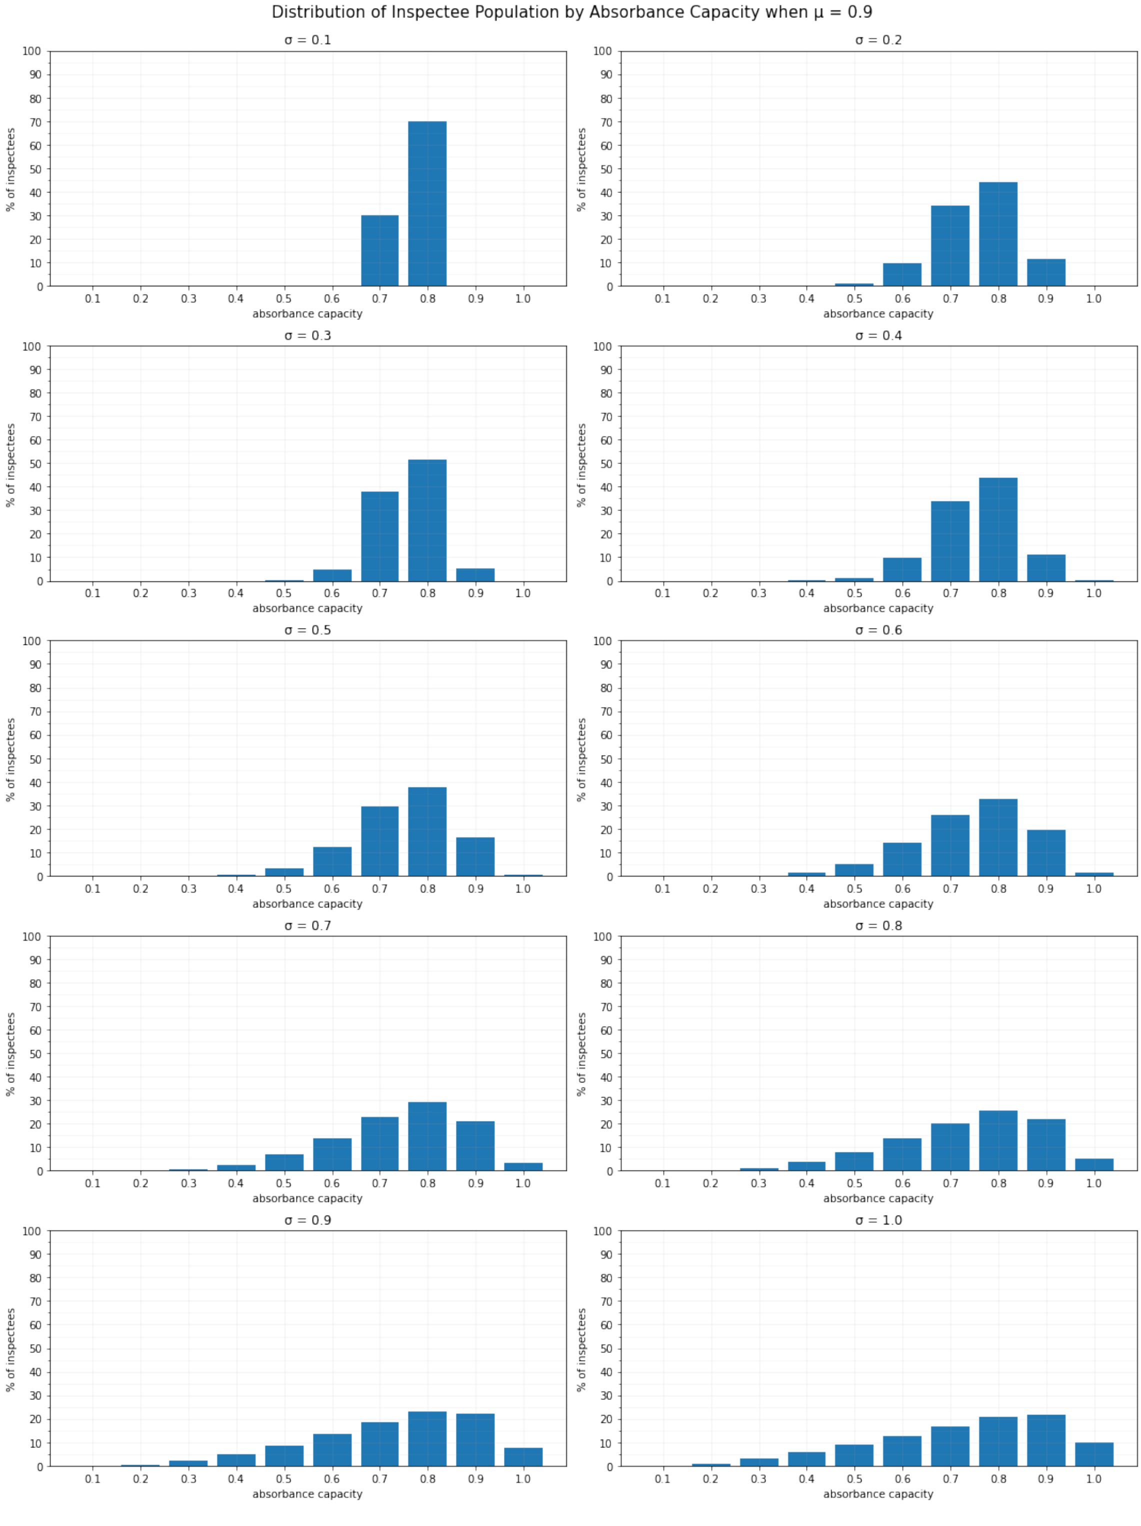

Figure A.38: Inspectee population distribution by absorbance capacity with μ=0.9.

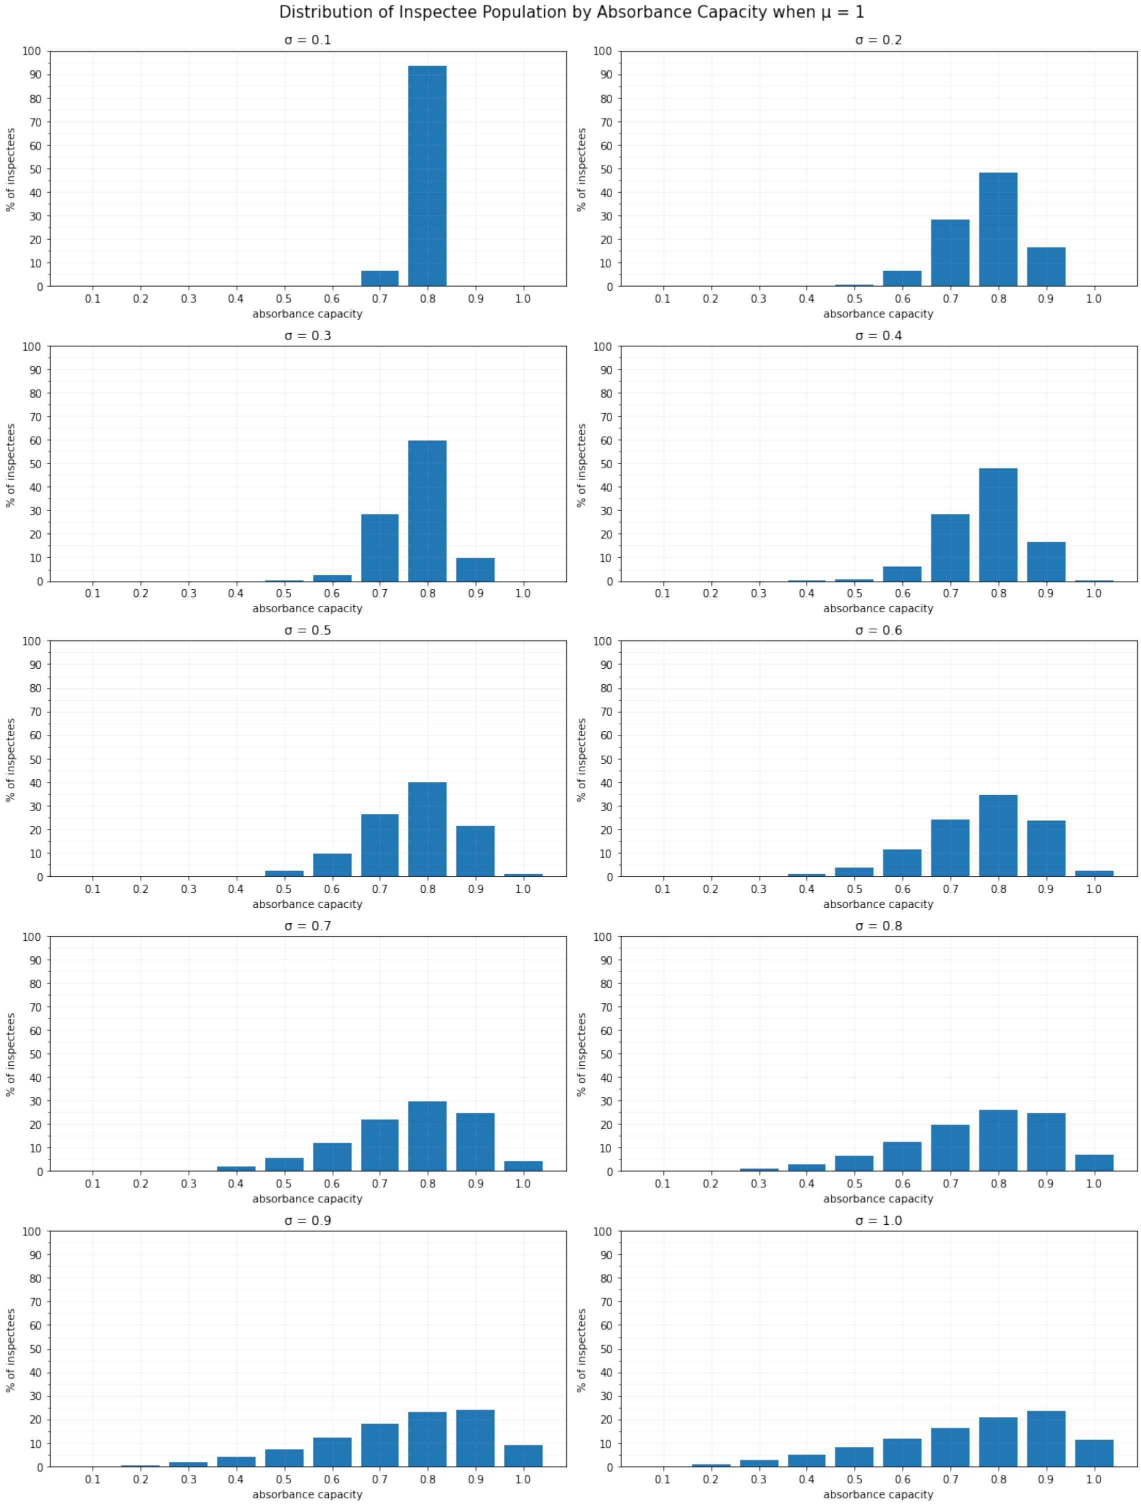

Figure A.39: Inspectee population distribution by absorbance capacity with μ=1.

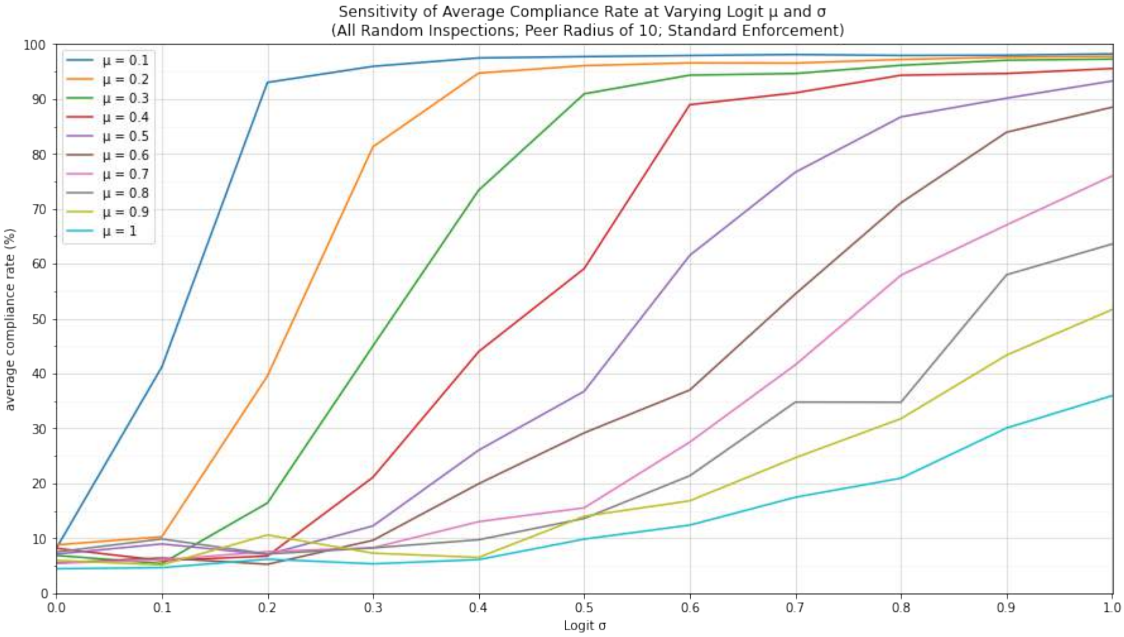

(a) Standard enforcement

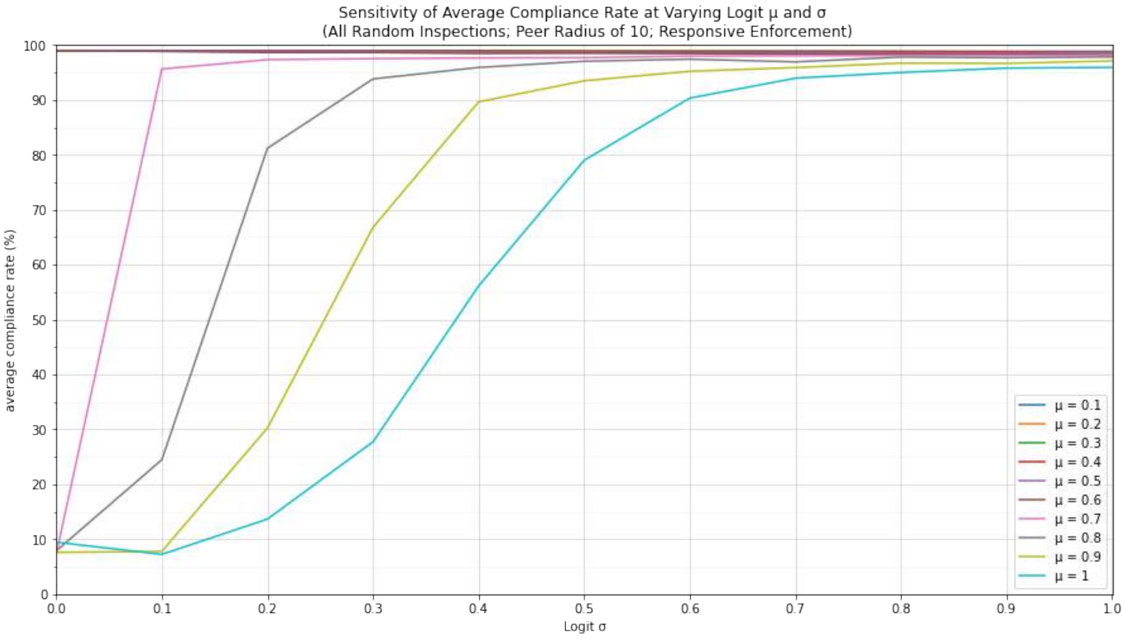

(b) Responsive enforcement.

Figure A.40: Sensitivity of average compliance rate with varying μ & σ (AR, r_{peers}=10).

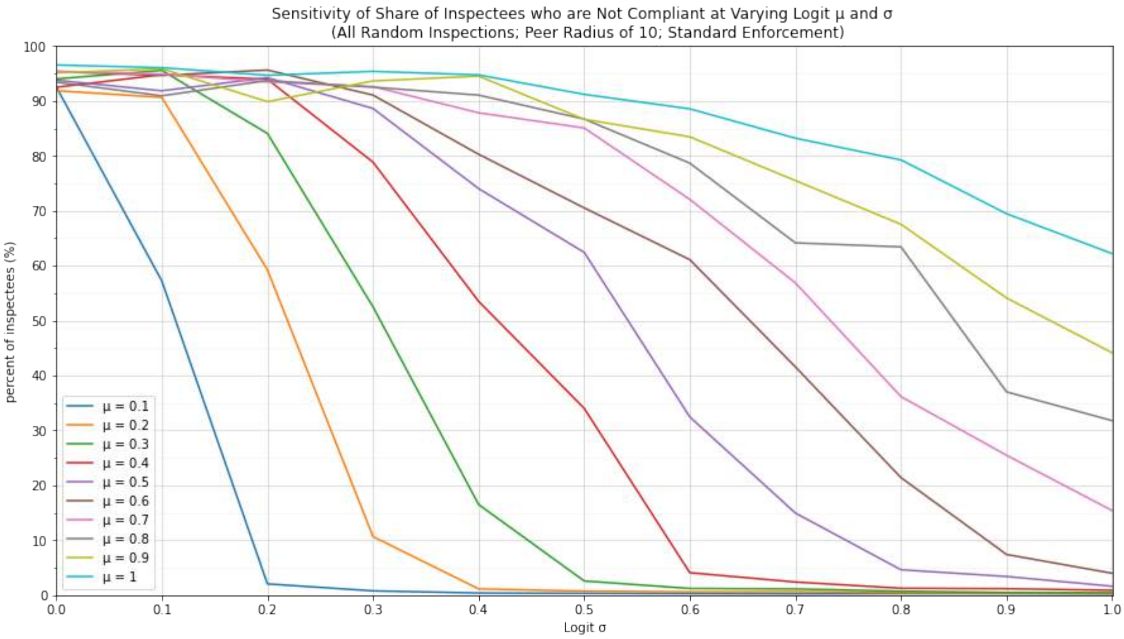

(a) Standard enforcement

For a detailed breakdown by compliance level, see Figure A.42.

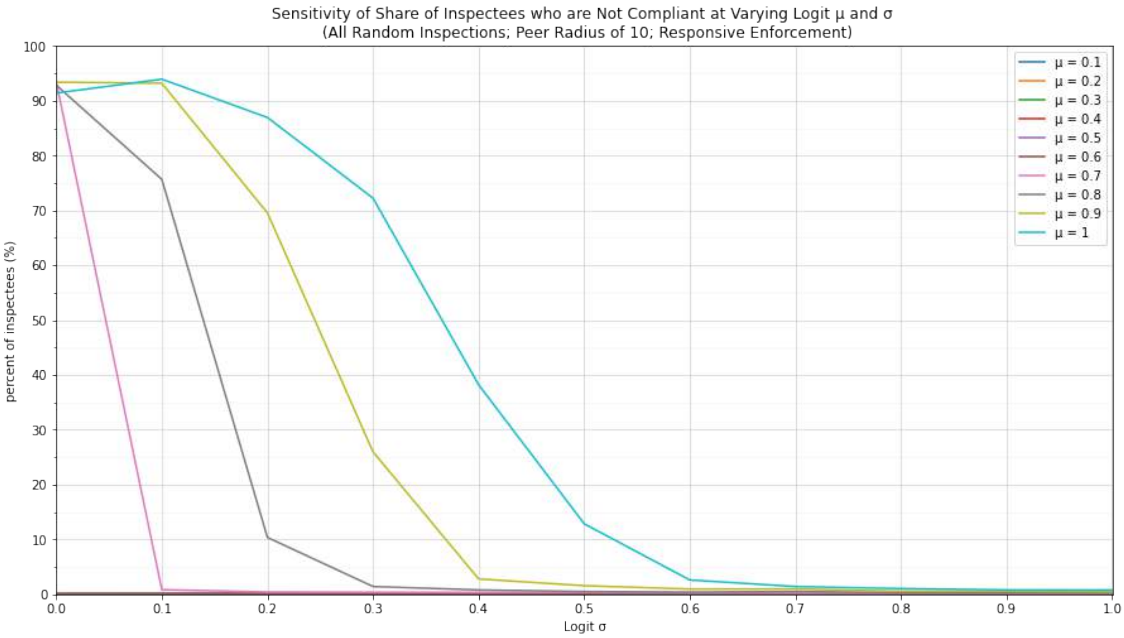

(b) Responsive enforcement.

For a detailed breakdown by compliance level, see Figure A.43.

Figure A.41: Sensitivity of the share of non-compliant inspectees with varying μ & σ (AR, r_{peers}=10).

Standard Enforcement

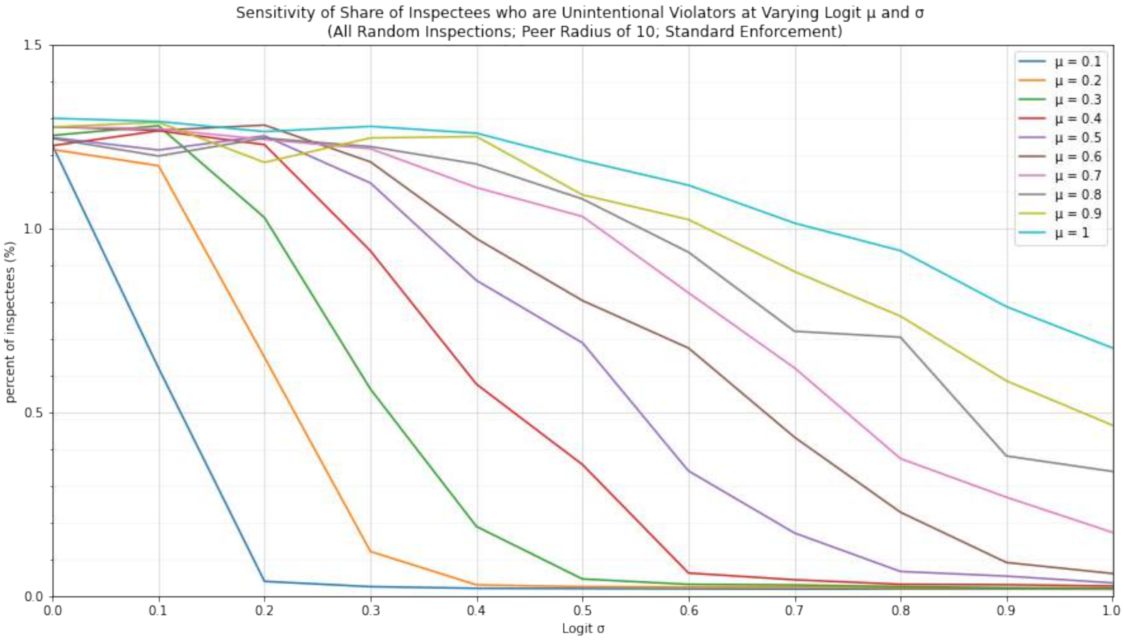

(a) Unintentional violators.

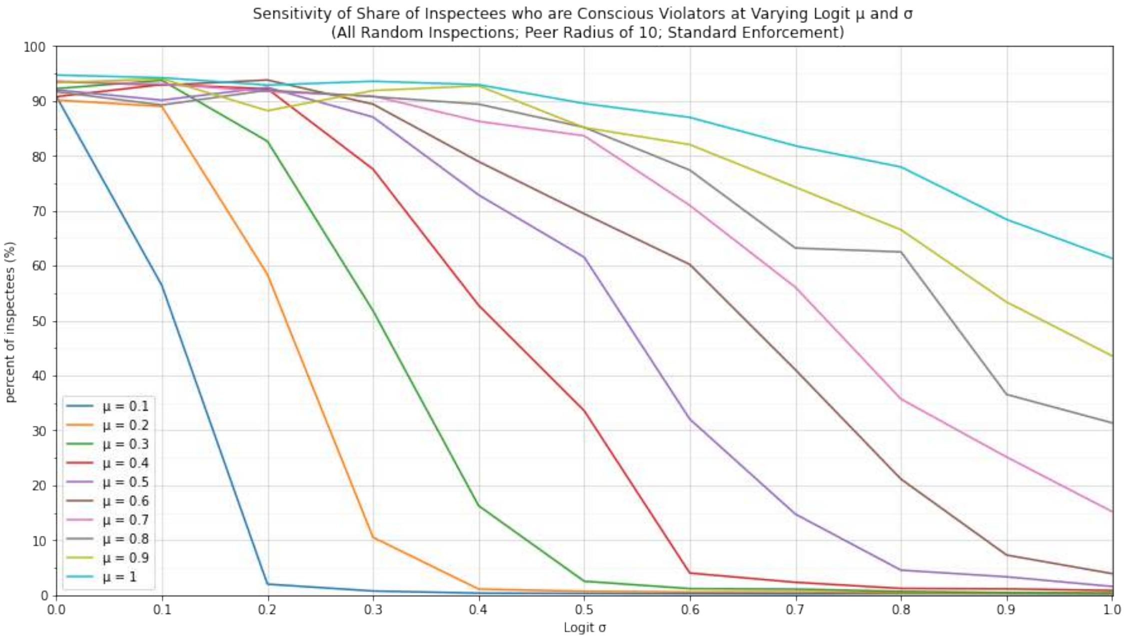

(b) Conscious violators.

Figure A.42: Sensitivity of the share of non-compliant inspectees by compliance level with varying μ & σ (AR, r_{peers}=10).

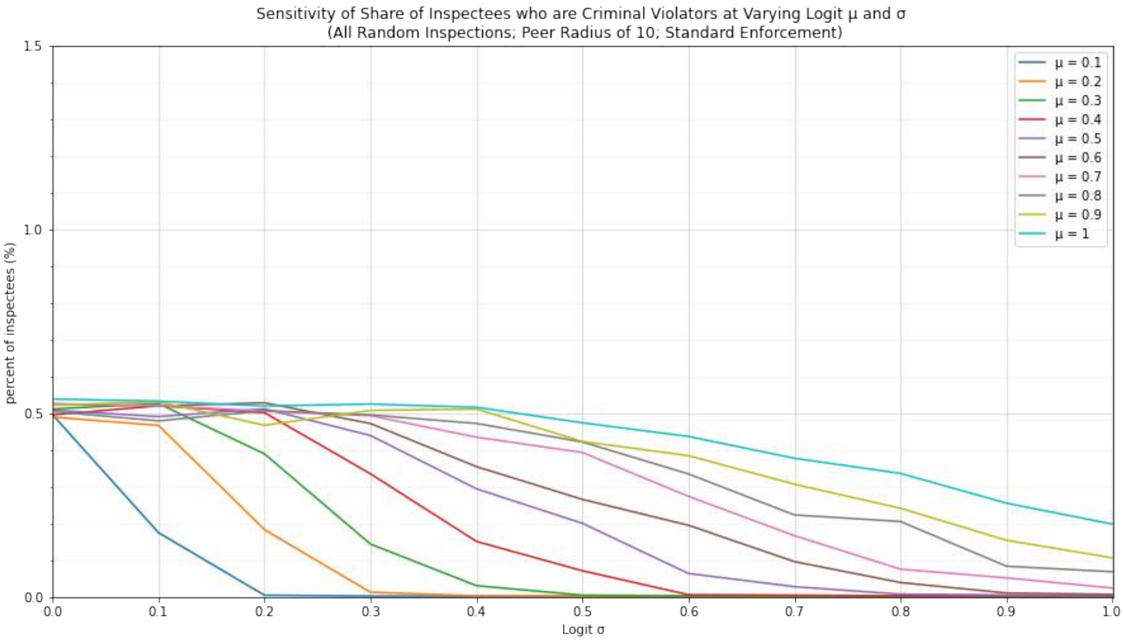

(c) Criminal violators.

Figure A.42: Sensitivity of the share of non-compliant inspectees by compliance level with varying μ & σ (AR, r_{peers}=10, continued).

Responsive Enforcement

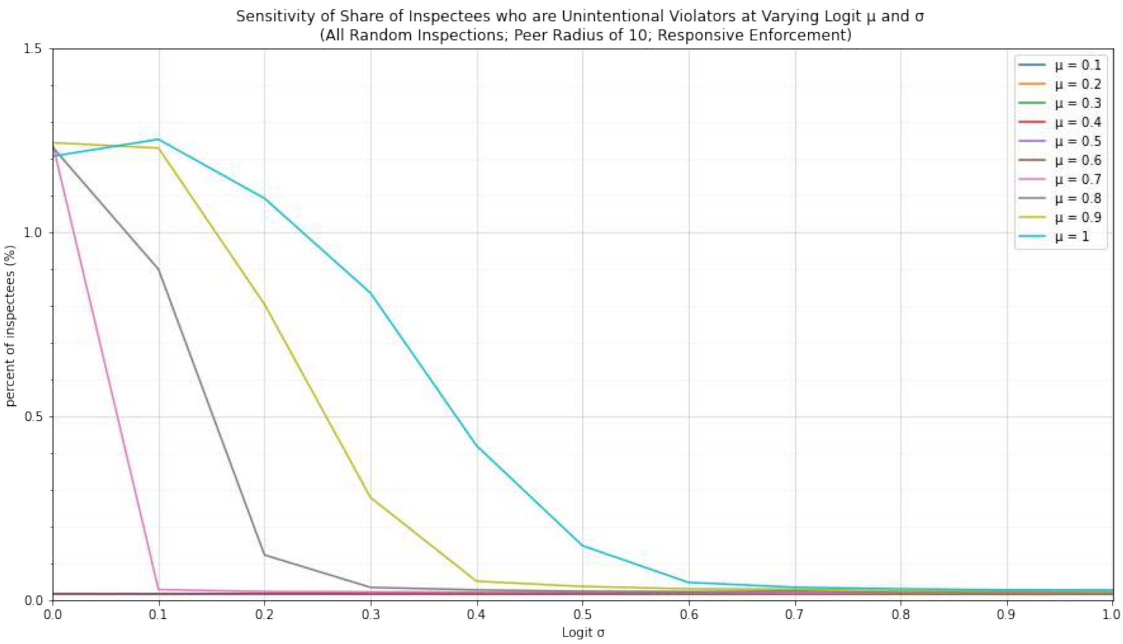

(a) Unintentional violators.

Figure A.43: Sensitivity of the share of non-compliant inspectees by compliance level with varying μ & σ (RE, AR, r_{peers}=10).

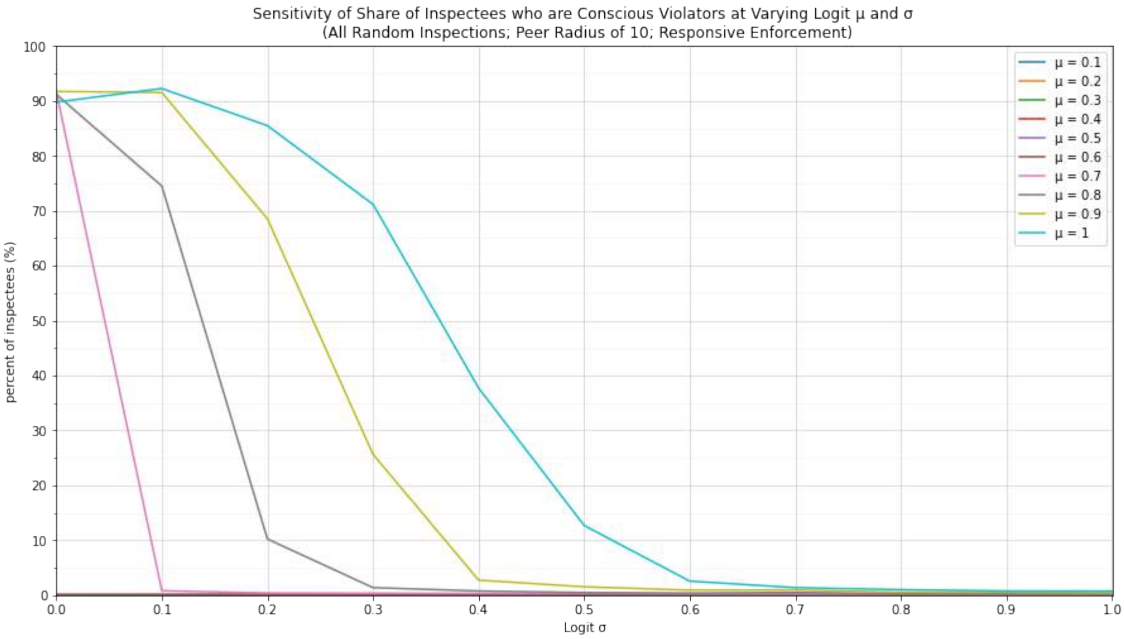

(b) Conscious violators.

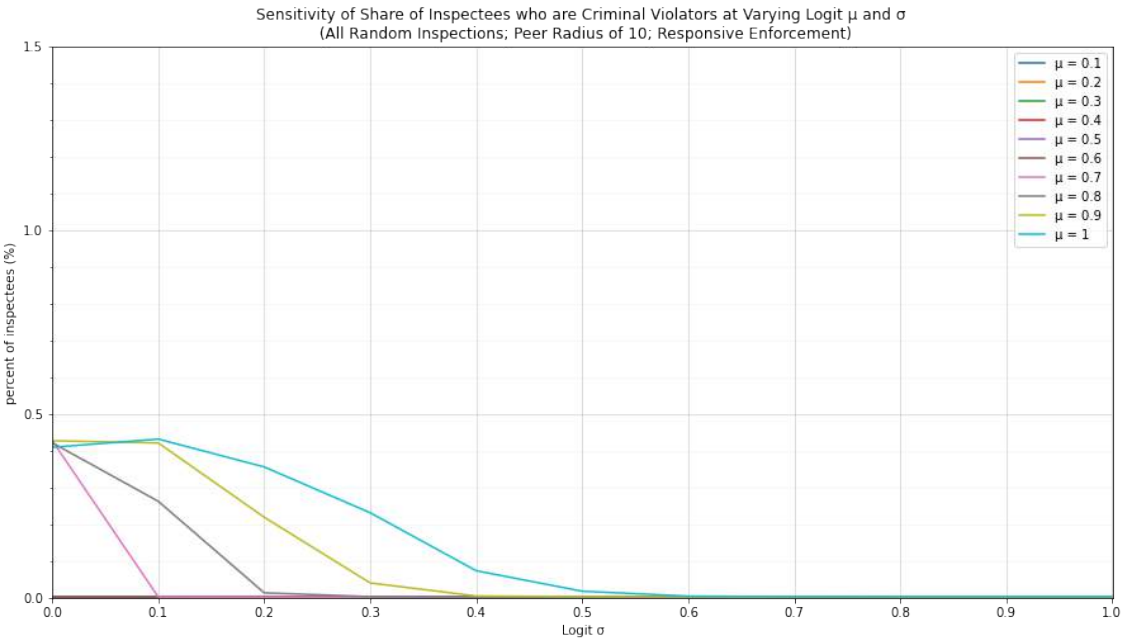

(c) Criminal violators.

Figure A.43: Sensitivity of the share of non-compliant inspectees by compliance level with varying μ & σ (RE, AR, r_{peers}=10, continued).

Varying Logit σ

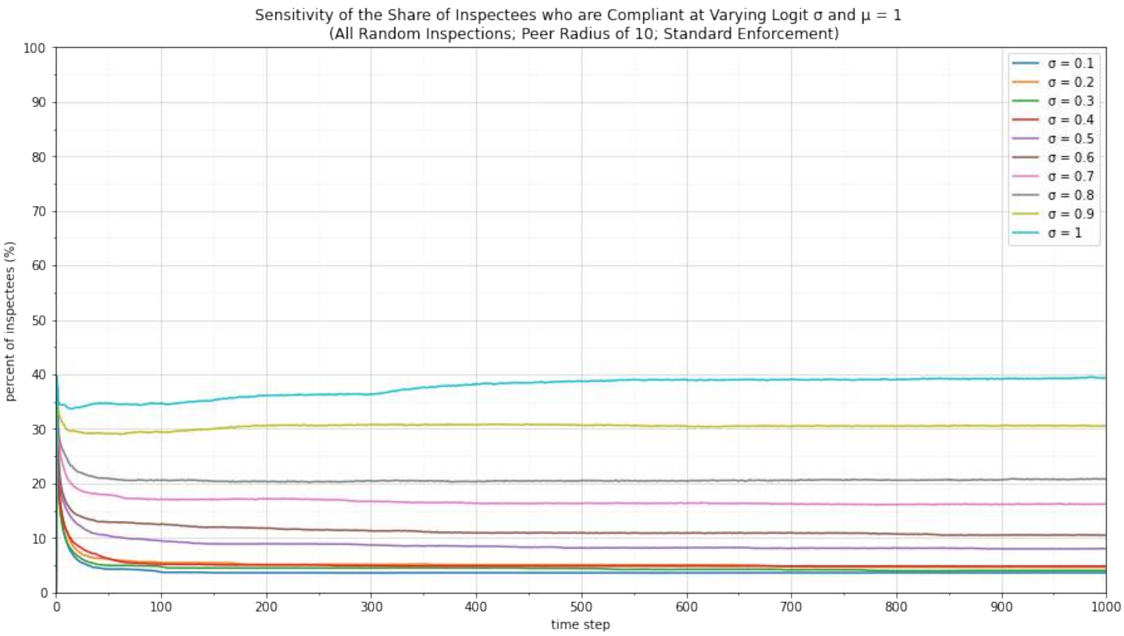

(a) Standard enforcement

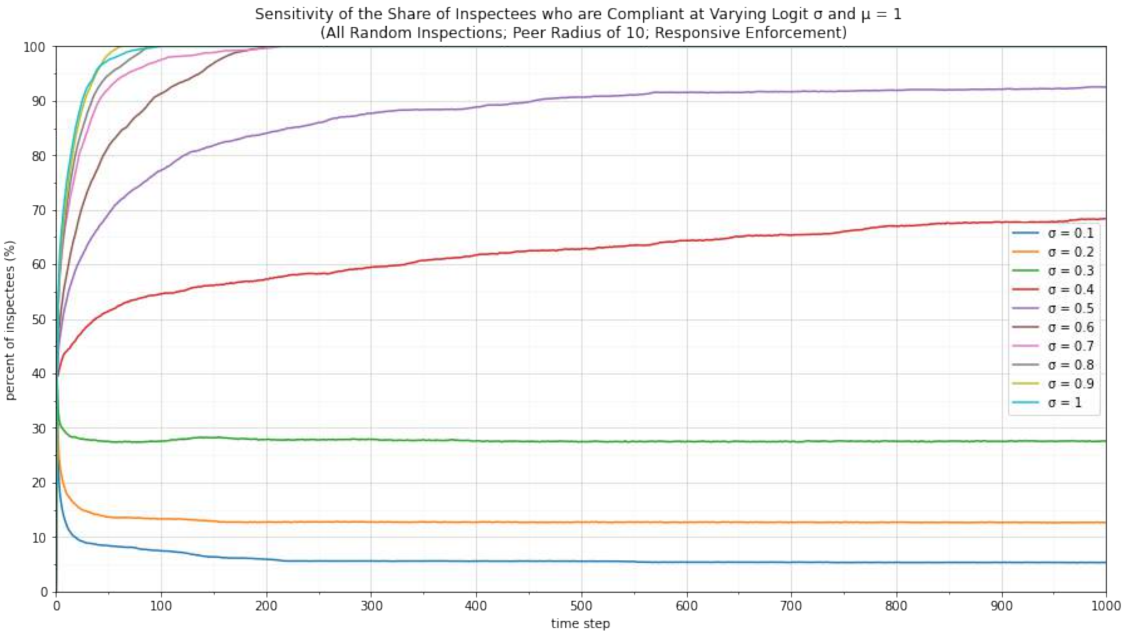

(b) Responsive enforcement.

Figure A.44: Share of compliant inspectees with varying σ & μ=1 (AR, r_{peers}=10).

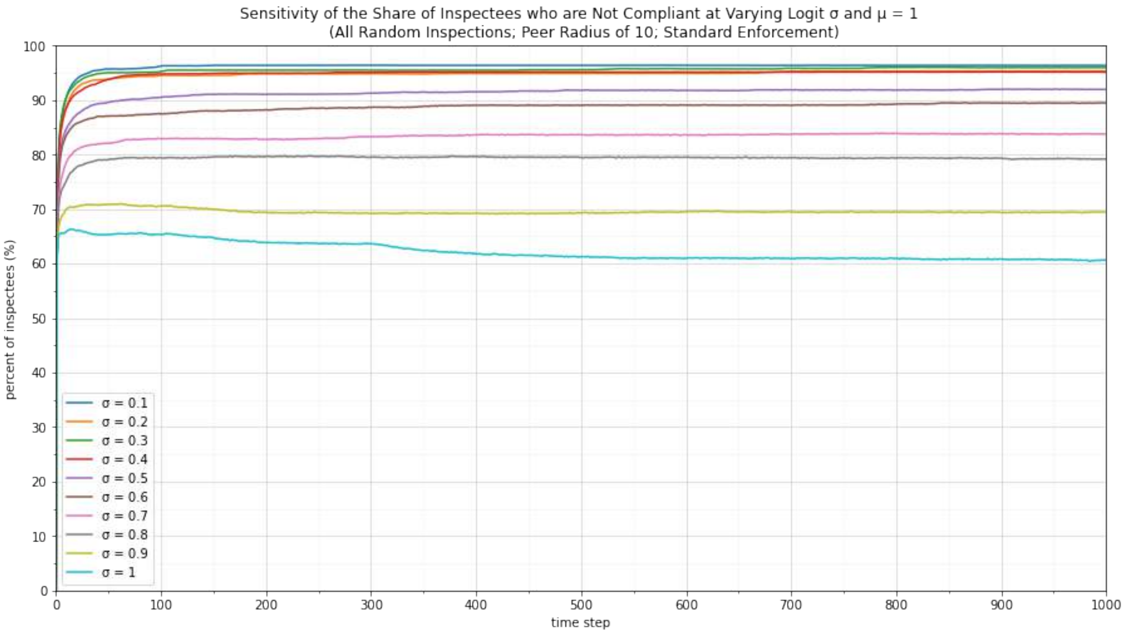

(a) Standard enforcement
For a detailed breakdown by compliance level, see Figure A.46.

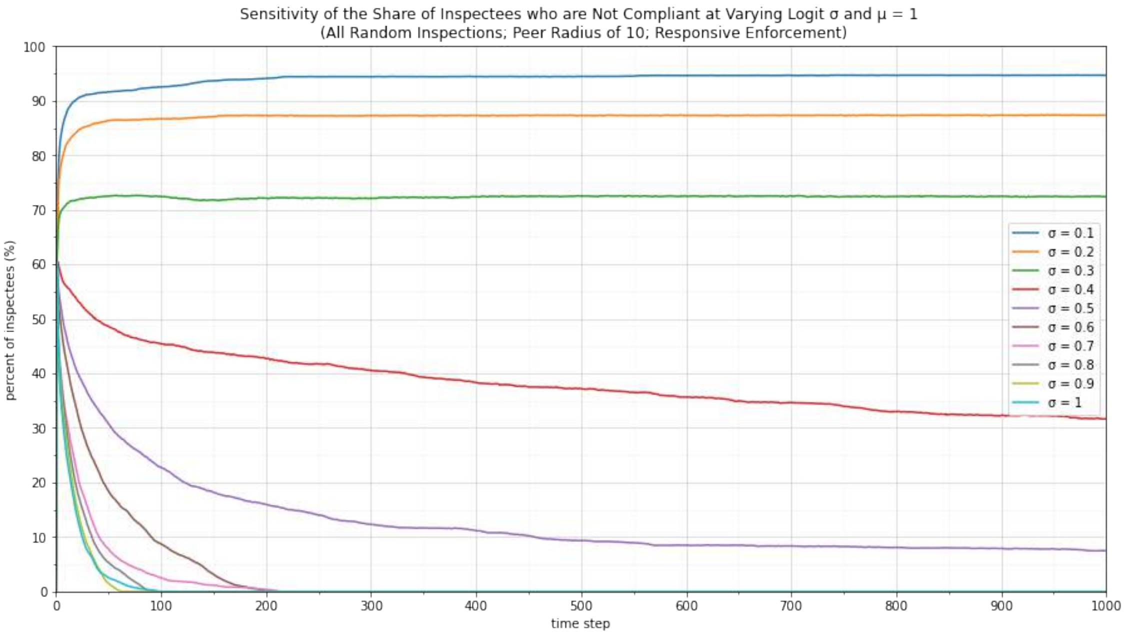

(b) Responsive enforcement.
For a detailed breakdown by compliance level, see Figure A.47.

Figure A.45: Share of non-compliant inspectees with varying σ & μ=1 (AR, r_{peers}=10).

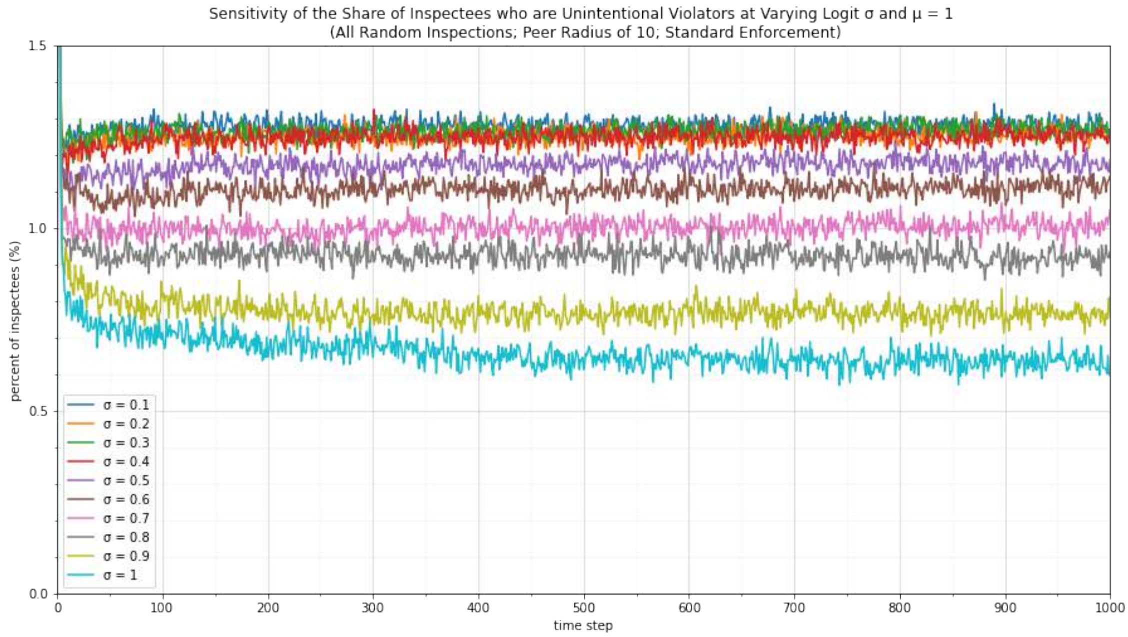

(a) Unintentional violators.

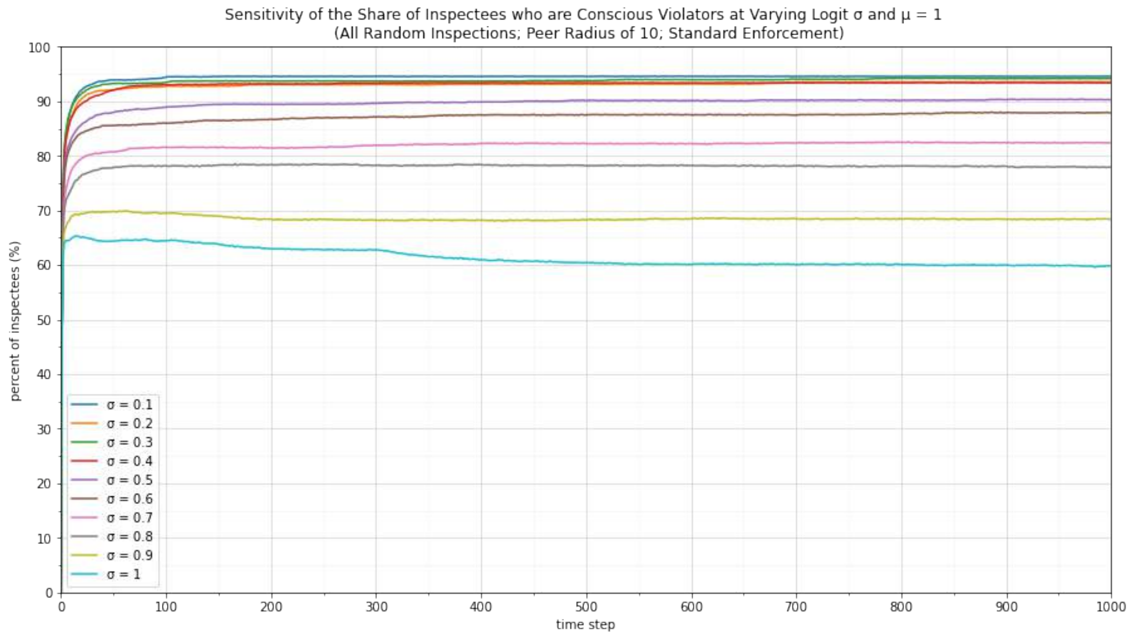

(b) Conscious violators.

Figure A.46: Non-compliant inspectee breakdown with varying σ & μ=1 (SE, AR, r_{peers}=10).

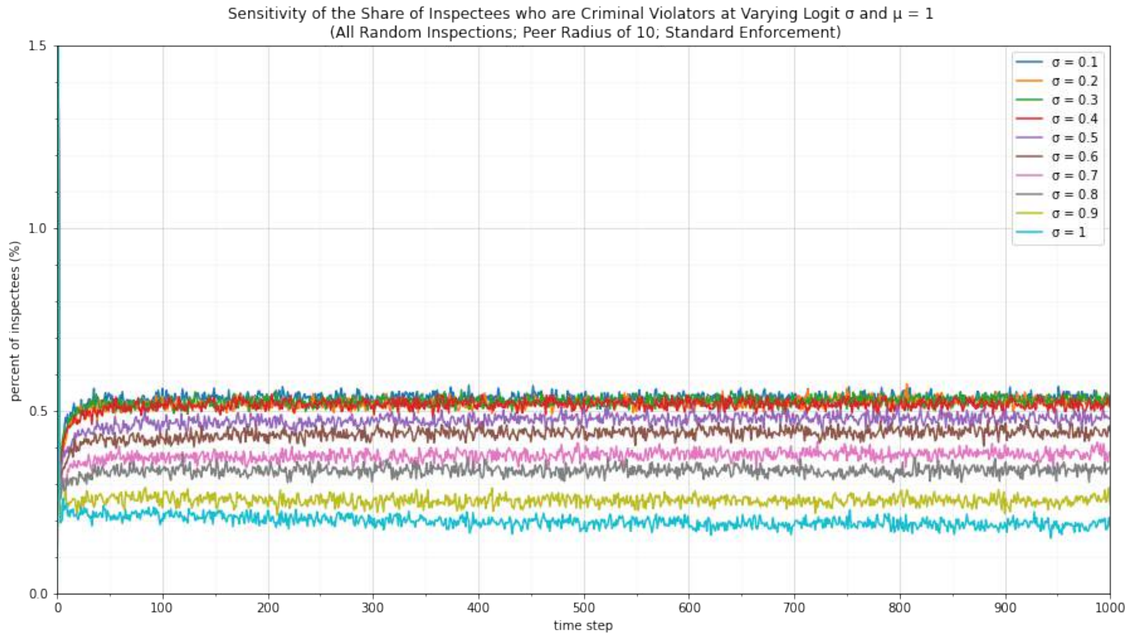

(c) Criminal violators.

Figure A.46: Non-compliant inspectee breakdown with varying σ & μ=1 (SE, AR, r_{peers}=10) (continued).

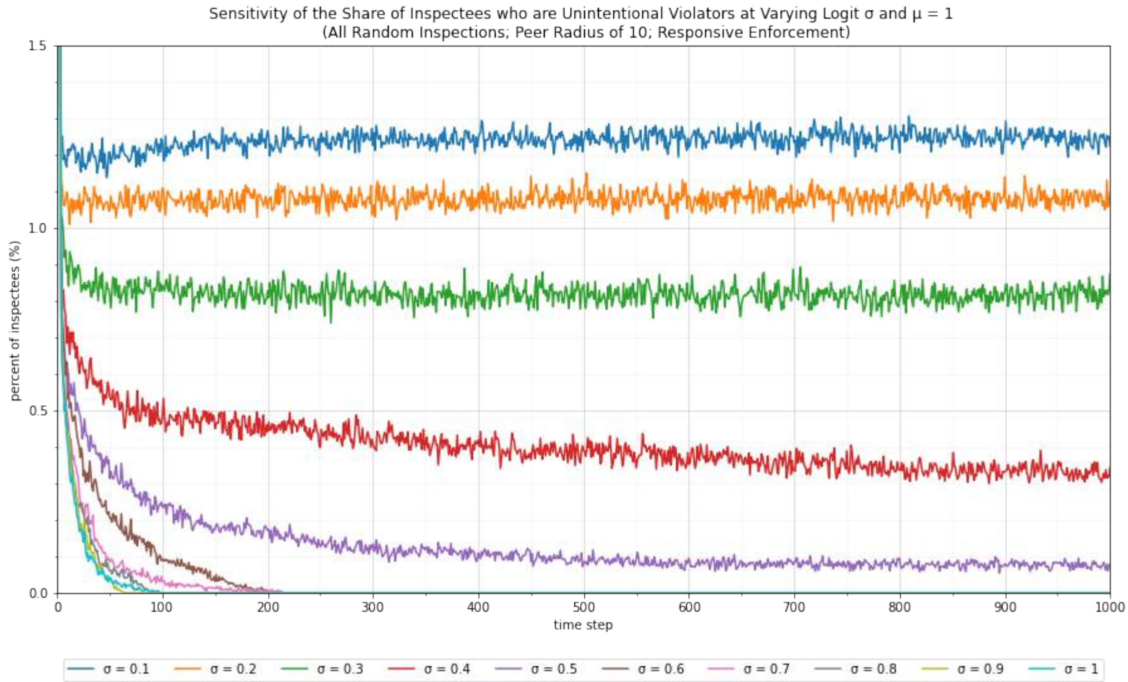

(a) Unintentional violators.

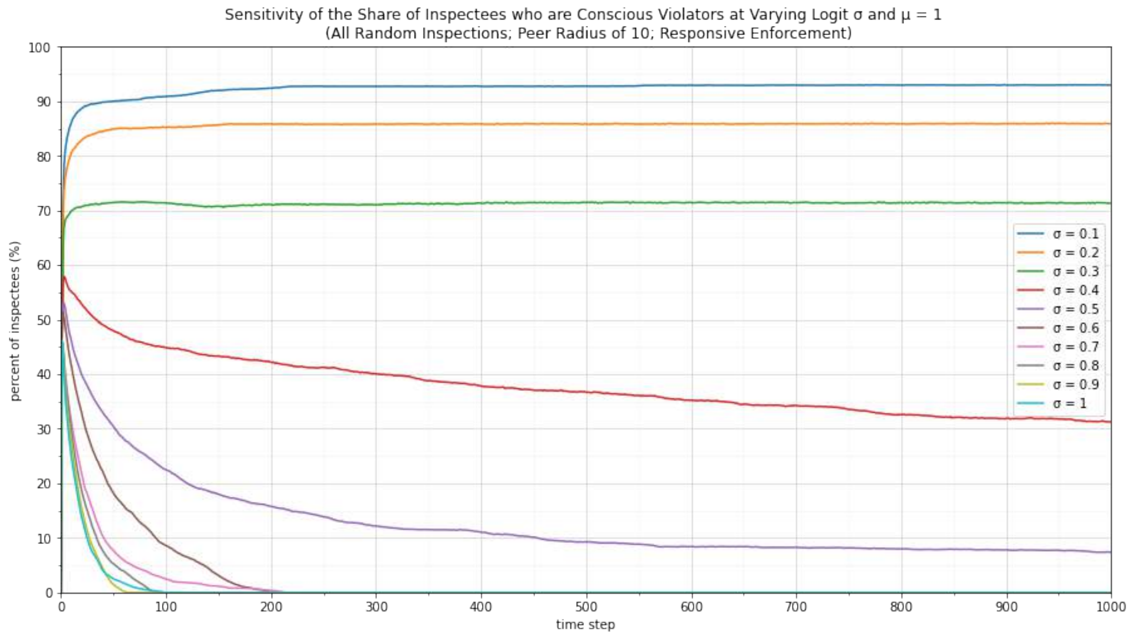

(b) Conscious violators.

Figure A.47: Non-compliant inspectee breakdown with varying σ & μ=1 (RE, AR, r_{peers}=10).

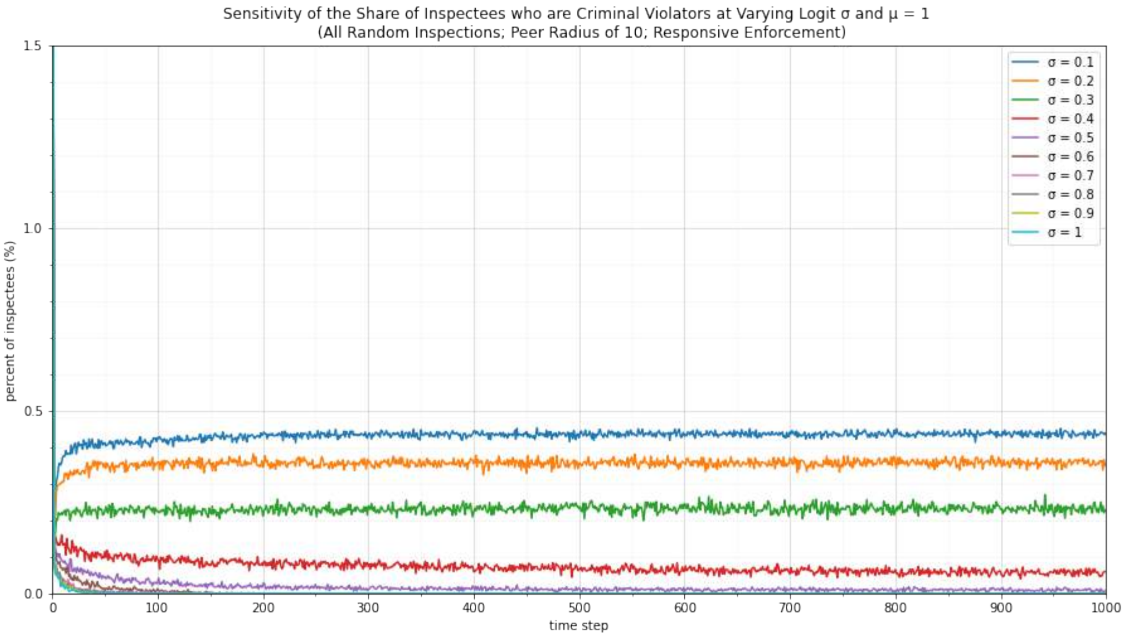

(c) Criminal violators.

Figure A.47: Non-compliant inspectee breakdown with varying σ & μ=1 (RE, AR, r_{peers}=10) (continued).

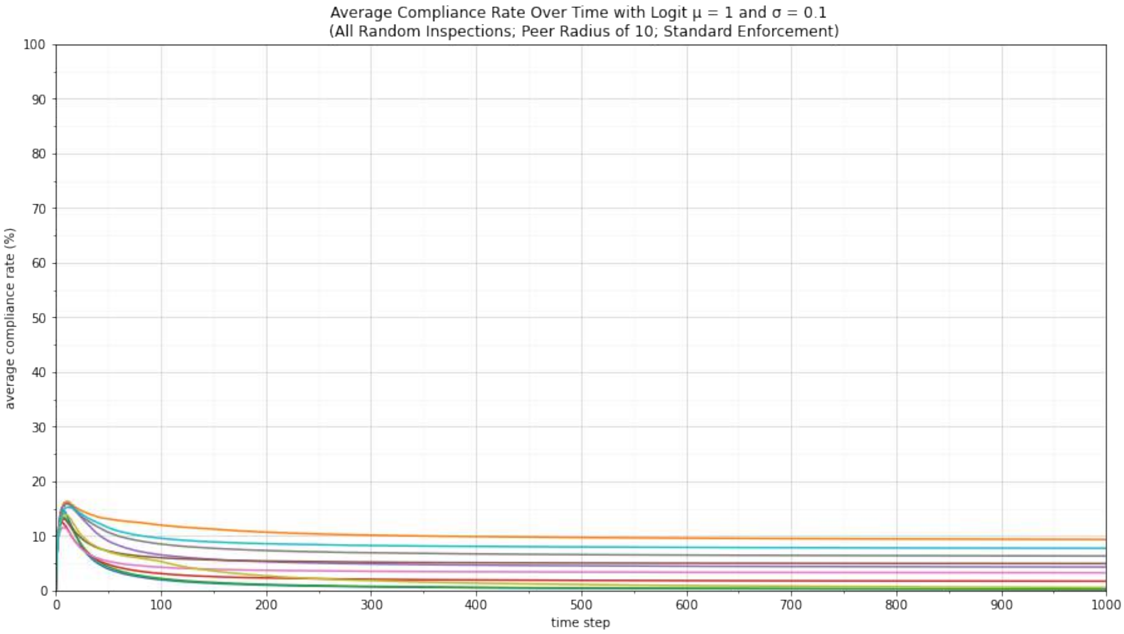

(a) σ = 0.1

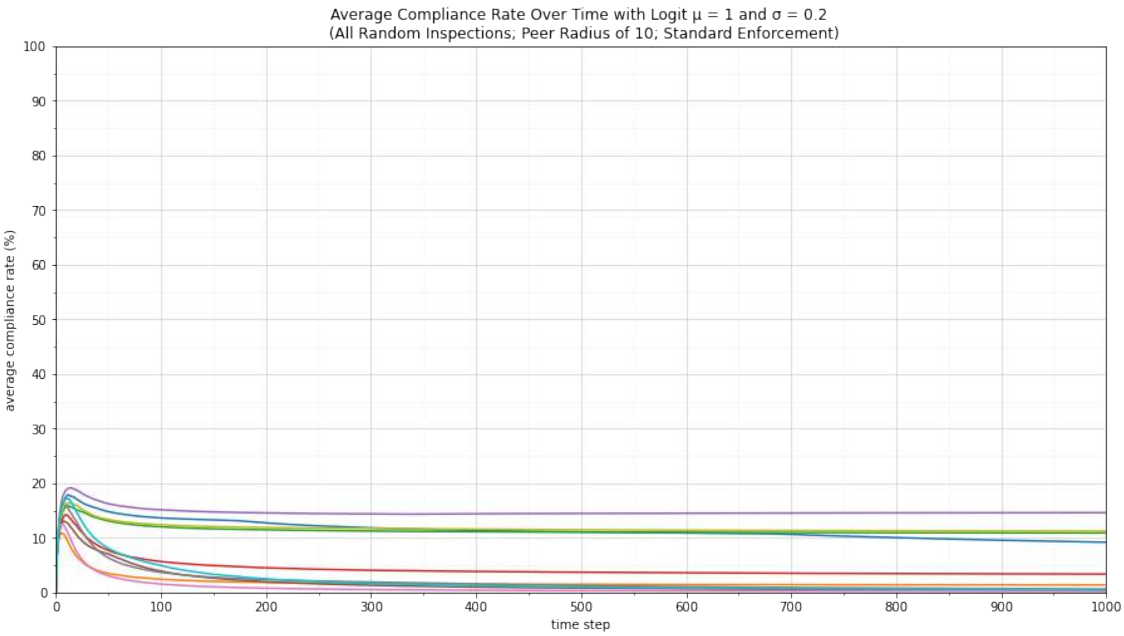

(b) σ = 0.2

Figure A.48: Average compliance rate with varying σ & μ=1 (SE, AR, r_{peers}=10).

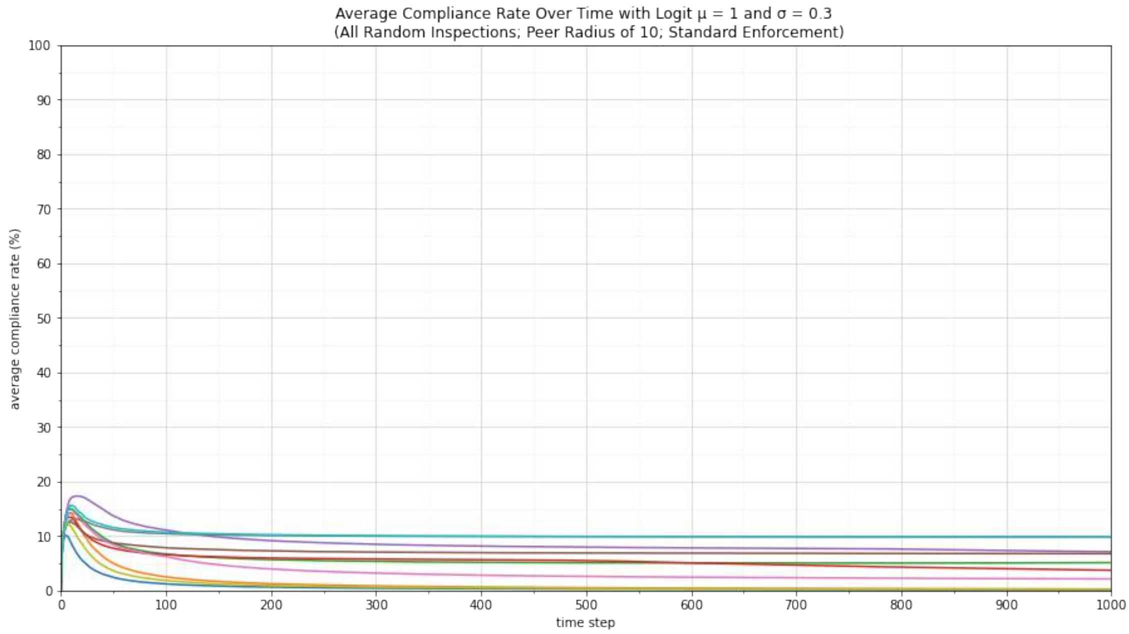

(c) $\sigma = 0.3$

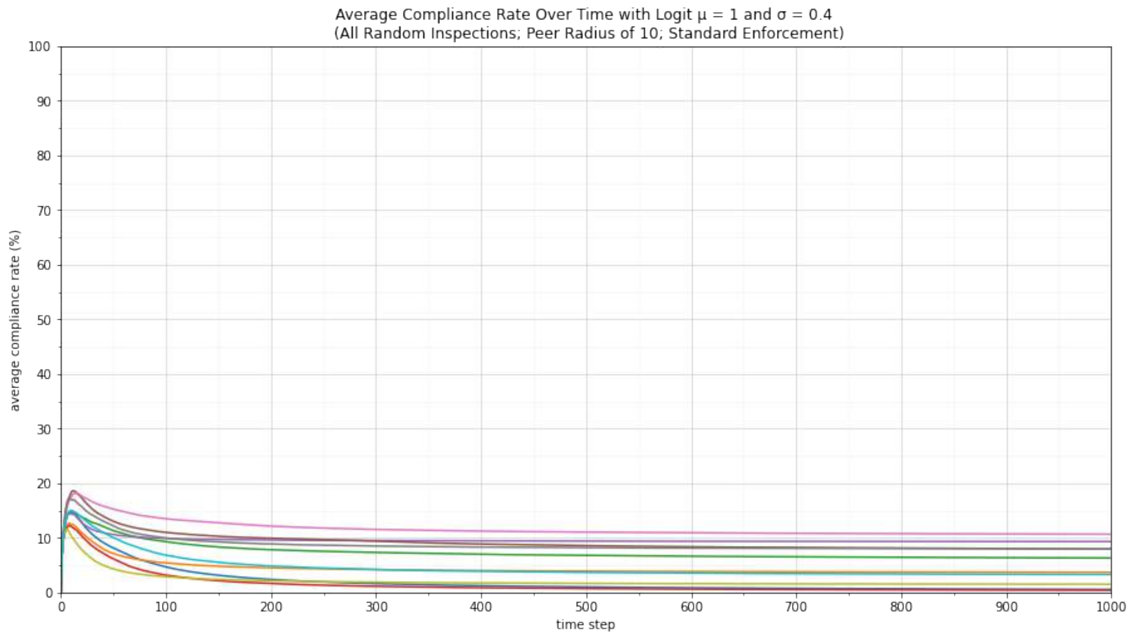

(d) $\sigma = 0.4$

Figure A.48: Average compliance rate with varying σ & μ=1 (SE, AR, r_{peers}=10, continued).

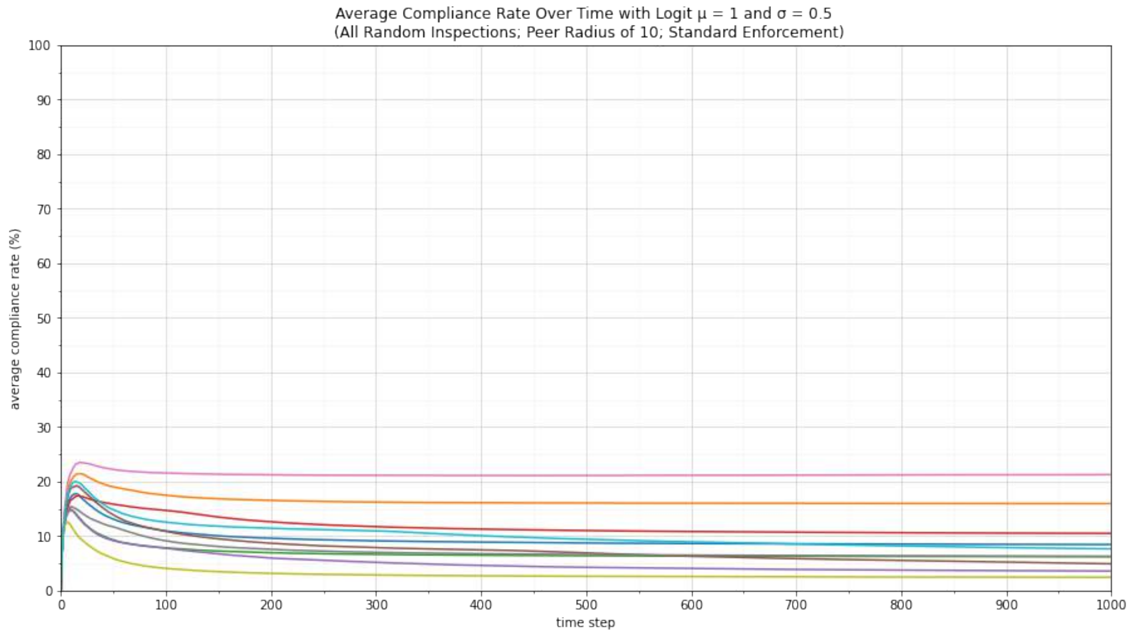

(e) σ = 0.5

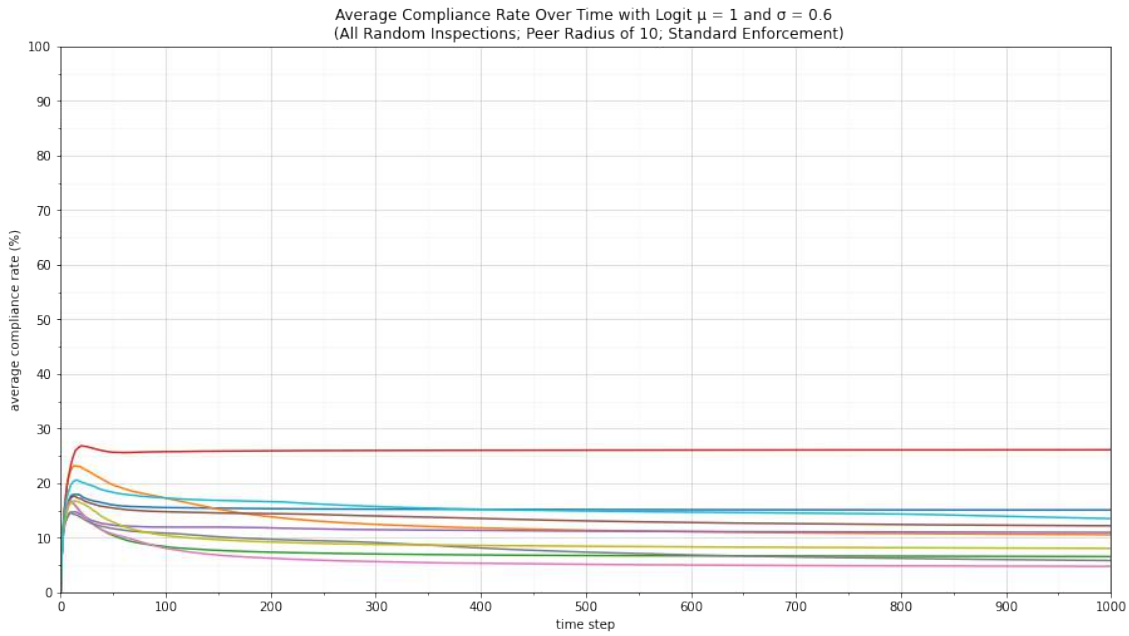

(f) σ = 0.6

Figure A.48: Average compliance rate with varying σ & μ=1 (SE, AR, r_{peers}=10, continued).

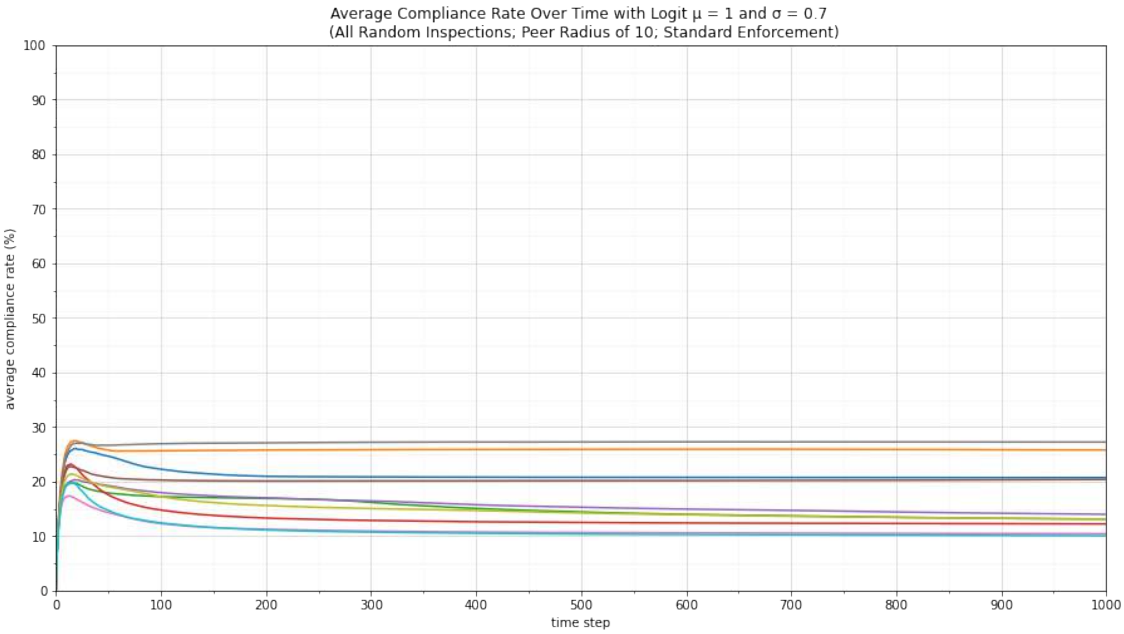

(g) $\sigma = 0.7$

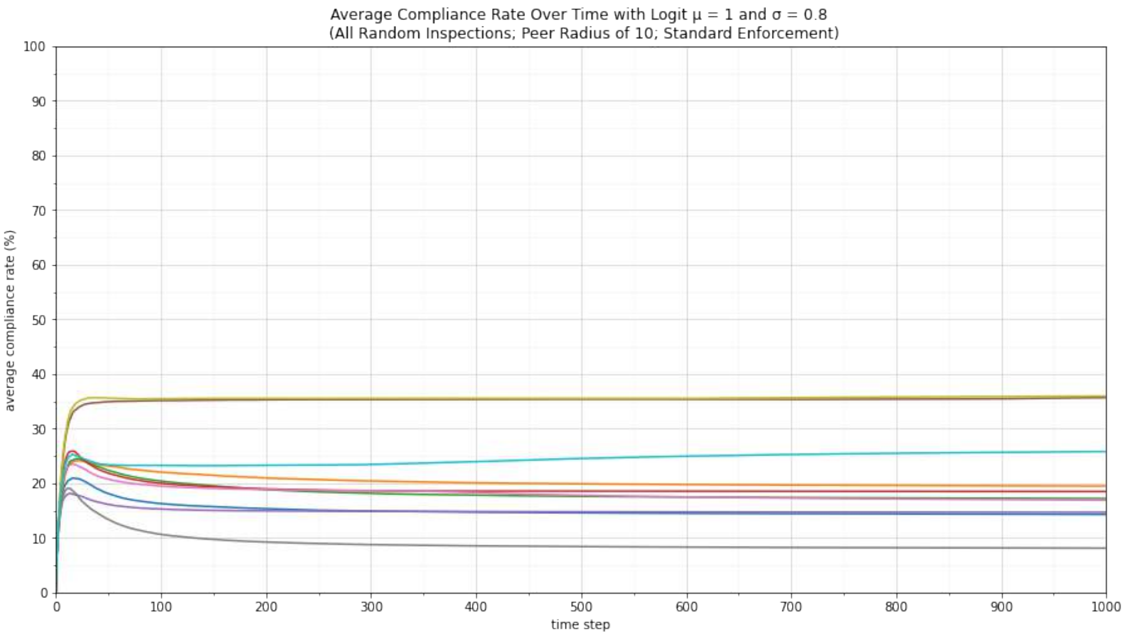

(h) $\sigma = 0.8$

Figure A.48: Average compliance rate with varying σ & μ=1 (SE, AR, r_{peers}=10, continued).

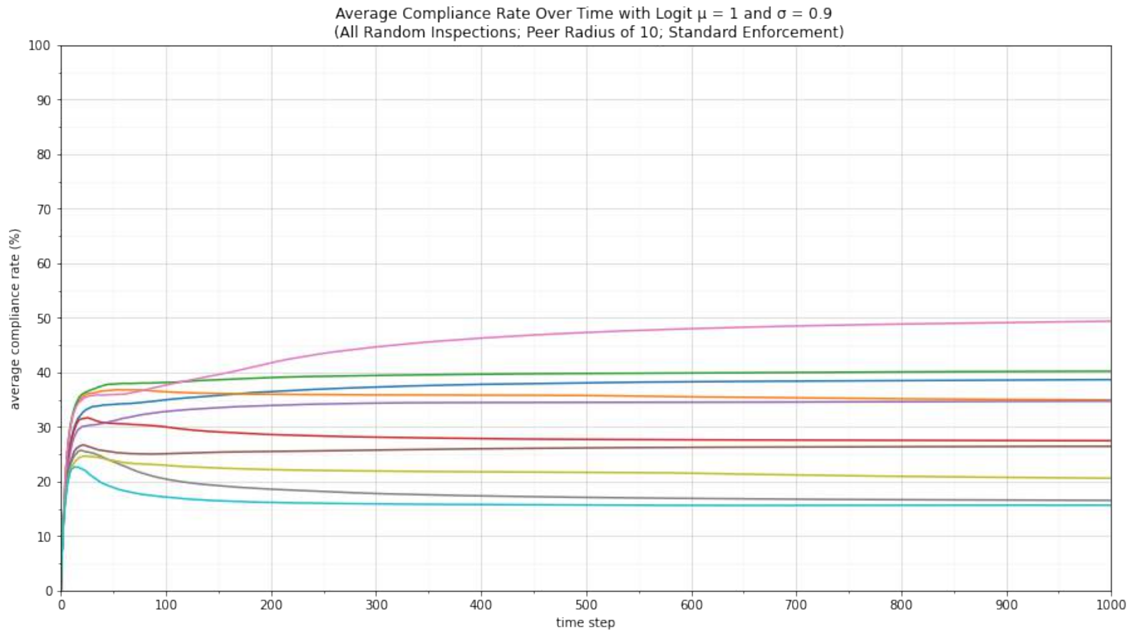

(i) $\sigma = 0.9$

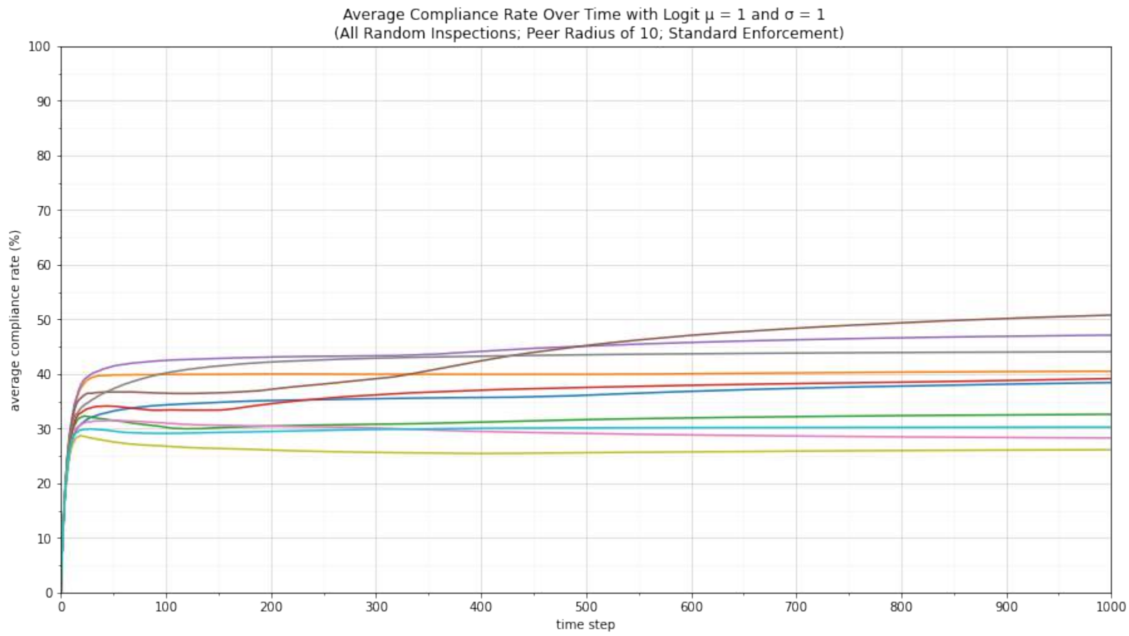

(j) $\sigma = 1$

Figure A.48: Average compliance rate with varying σ & μ=1 (SE, AR, r_{peers}=10, continued).

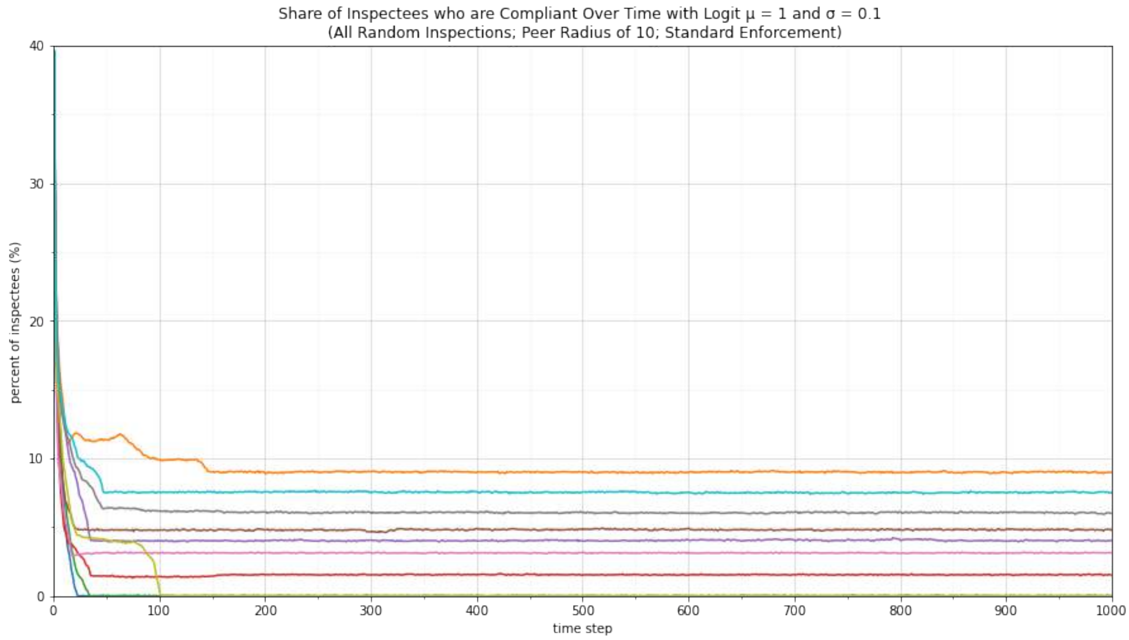

(a) $\sigma = 0.1$

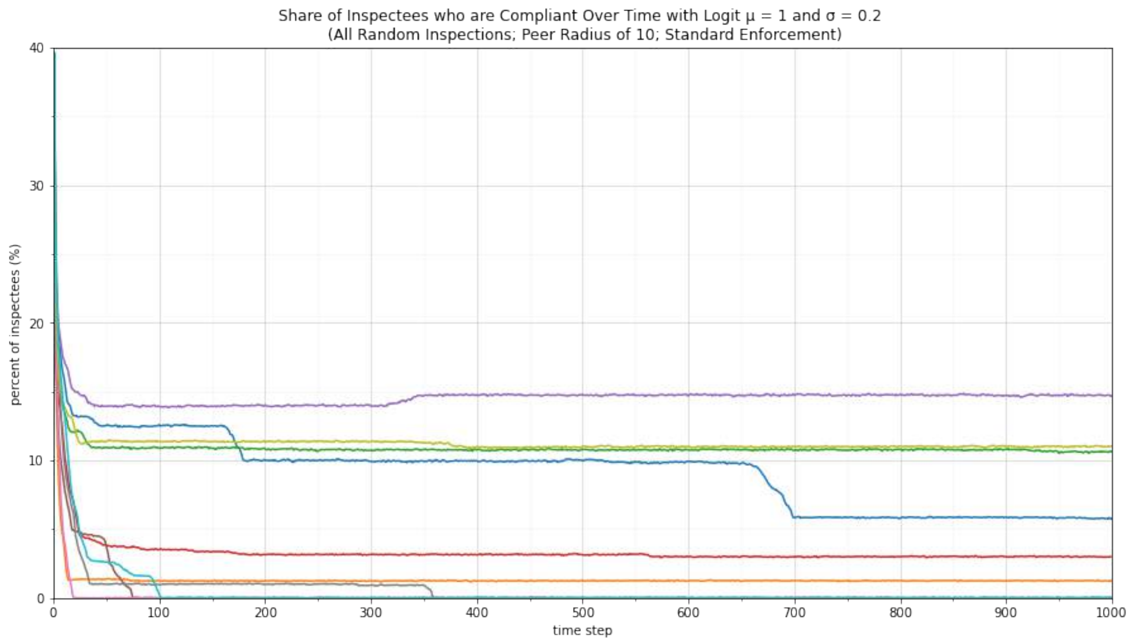

(b) $\sigma = 0.2$

Figure A.49: Share of compliant inspectees with varying σ & μ=1 (SE, AR, r_{peers}=10).

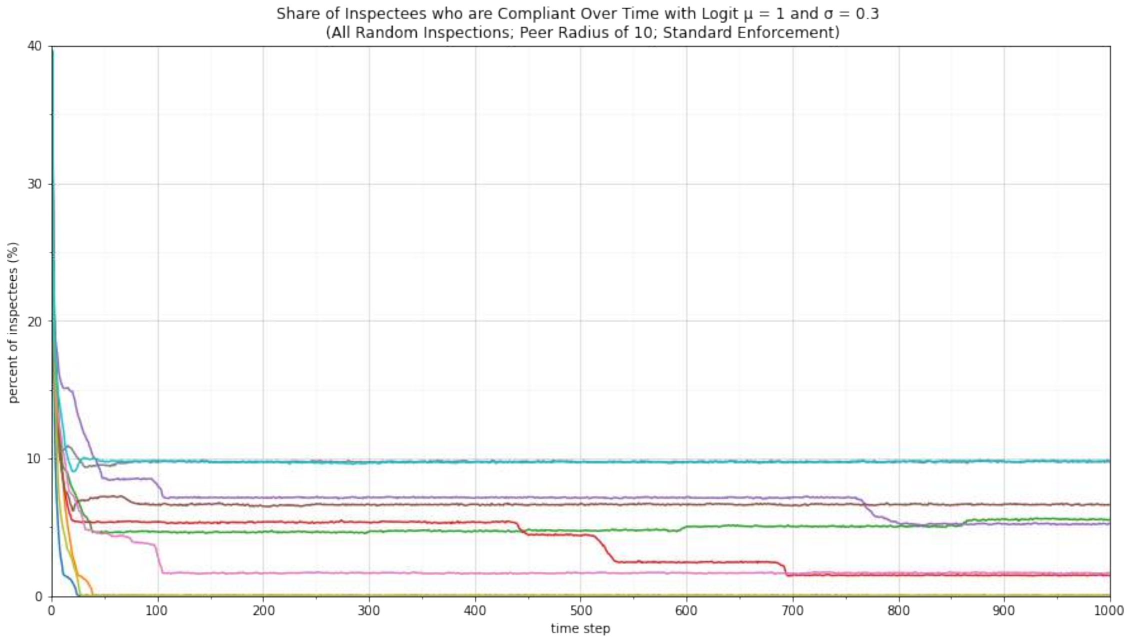

(c) σ = 0.3

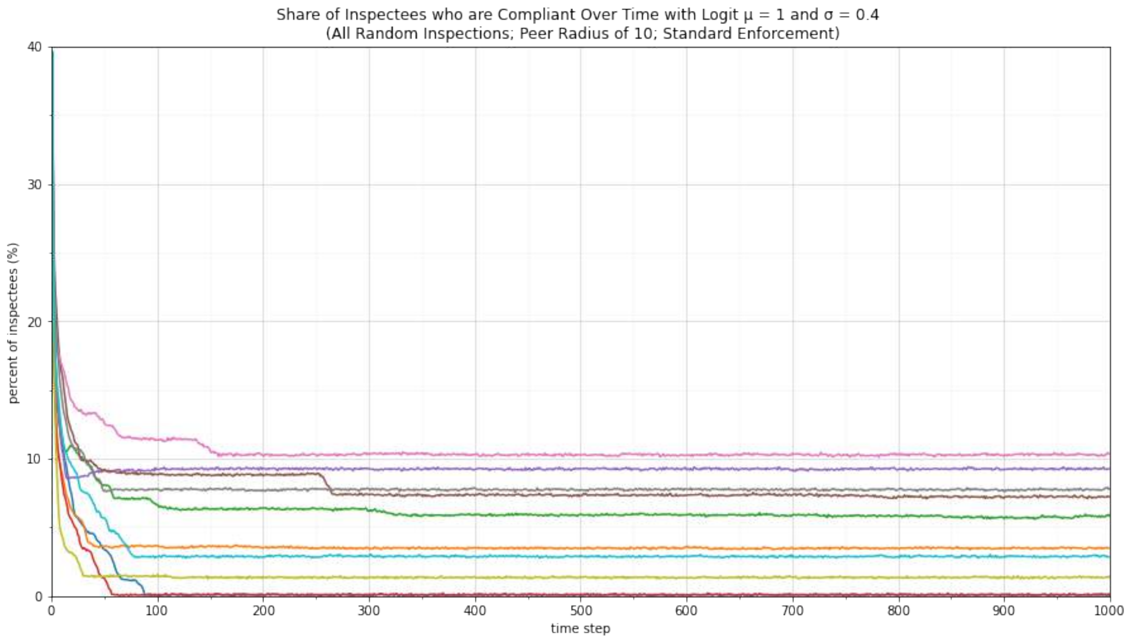

(d) σ = 0.4

Figure A.49: Share of compliant inspectees with varying σ & μ=1 (SE, AR, r_{peers}=10, continued).

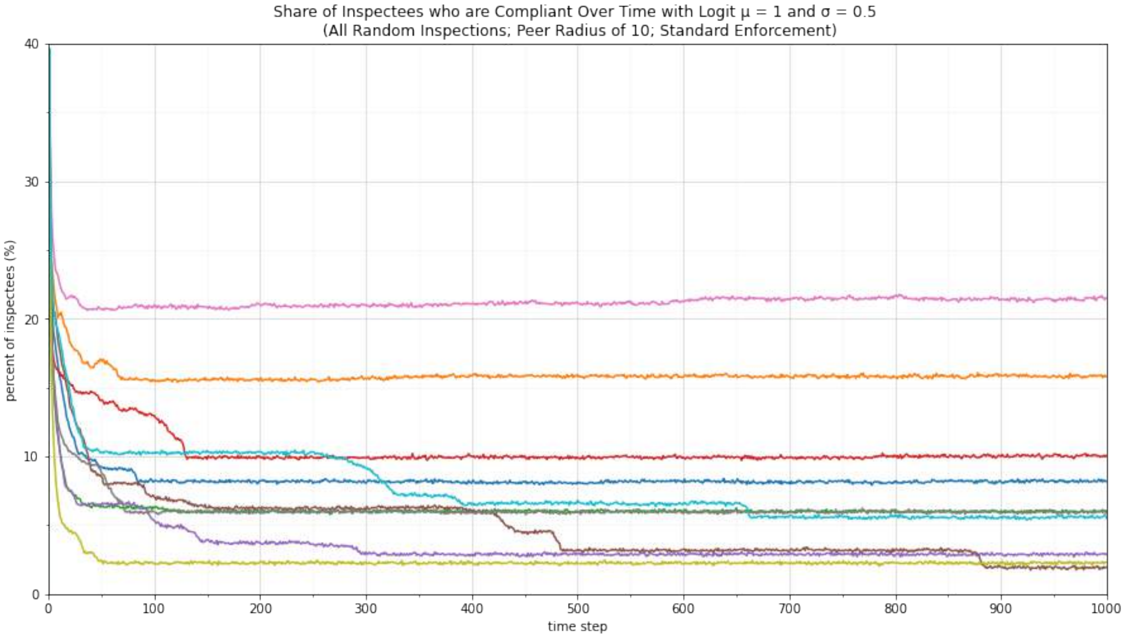

(e) σ = 0.5

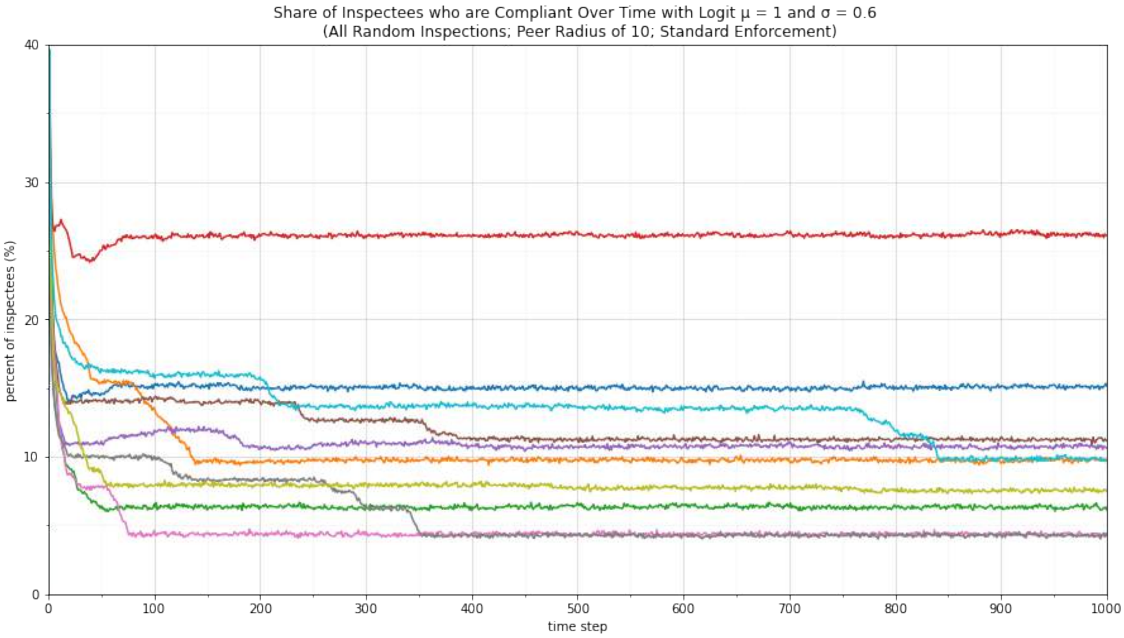

(f) σ = 0.6

Figure A.49: Share of compliant inspectees with varying σ & μ=1 (SE, AR, r_{peers}=10, continued).

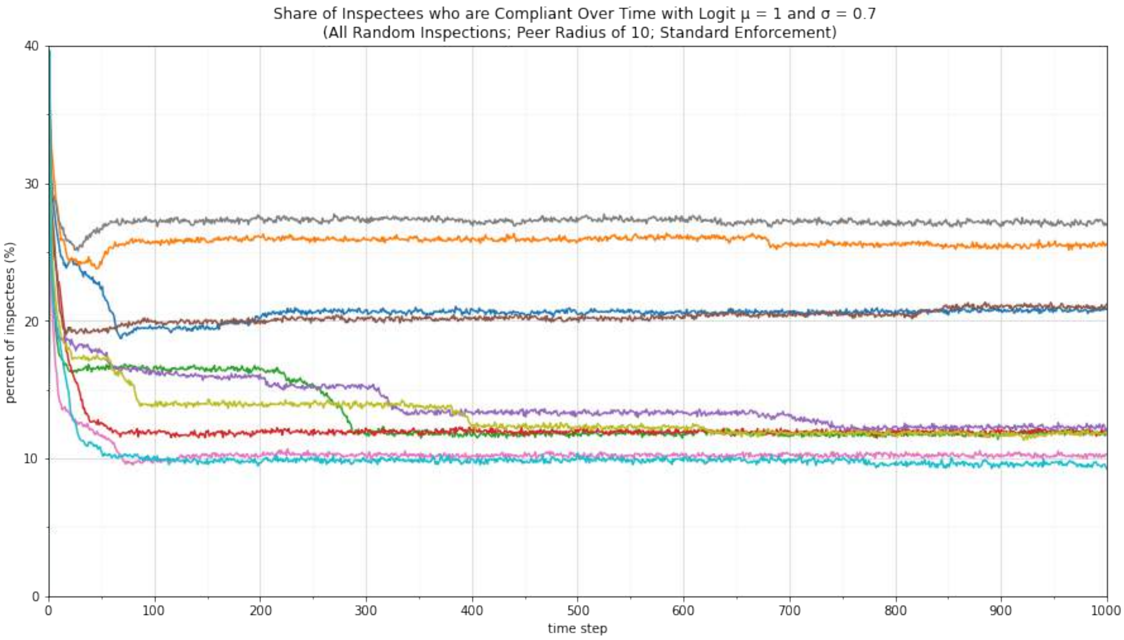

(g) σ = 0.7

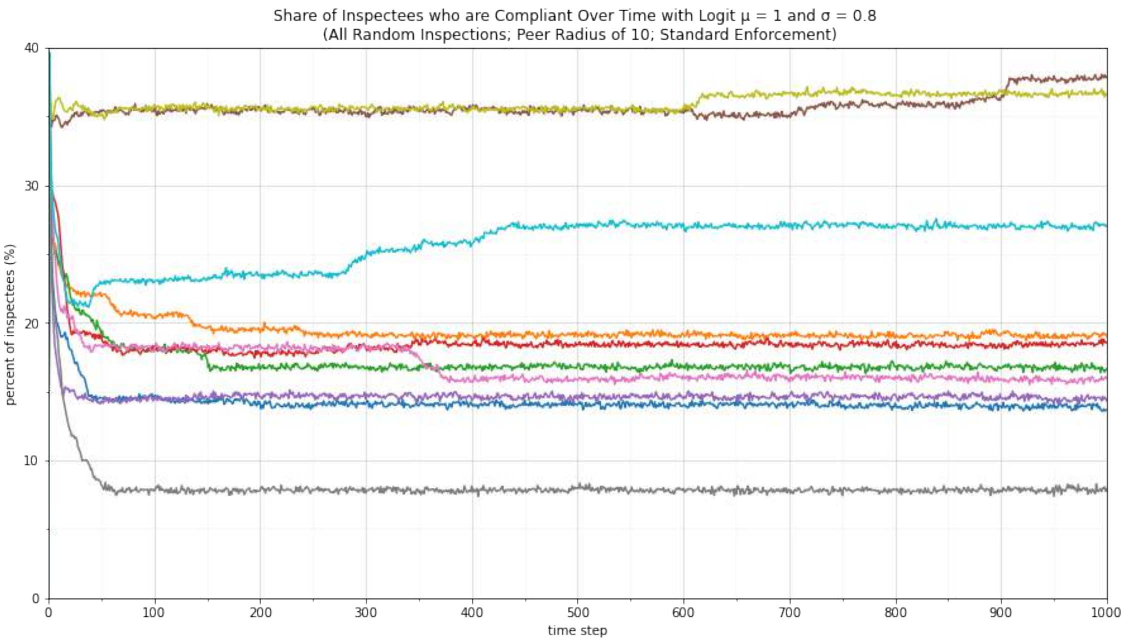

(h) σ = 0.8

Figure A.49: Share of compliant inspectees with varying σ & μ=1 (SE, AR, r_{peers}=10, continued).

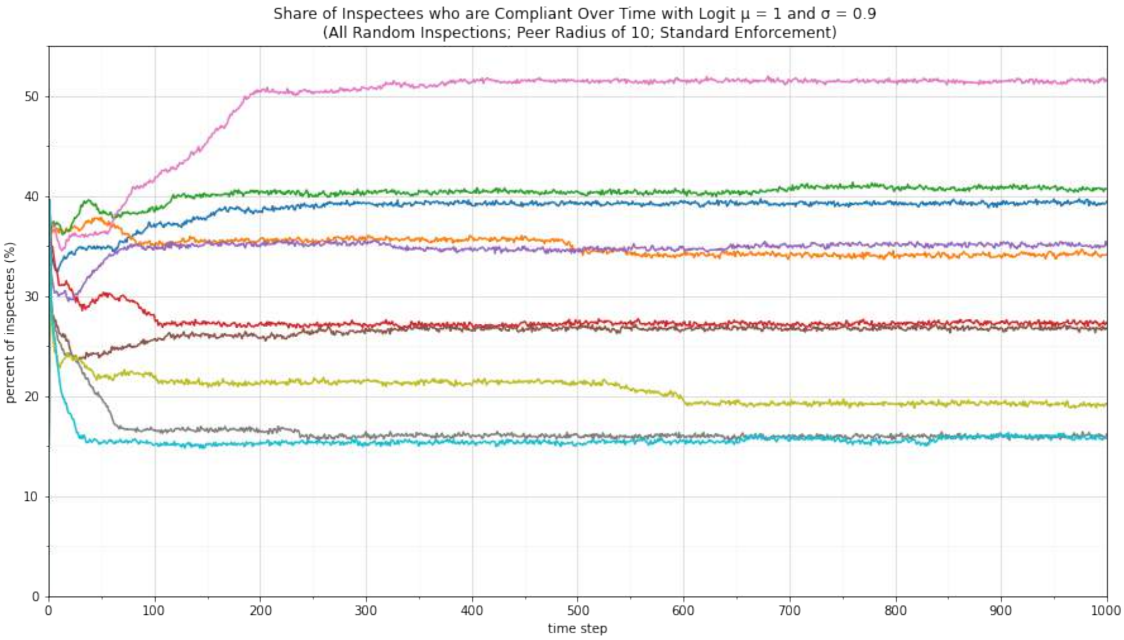

(i) $\sigma = 0.9$

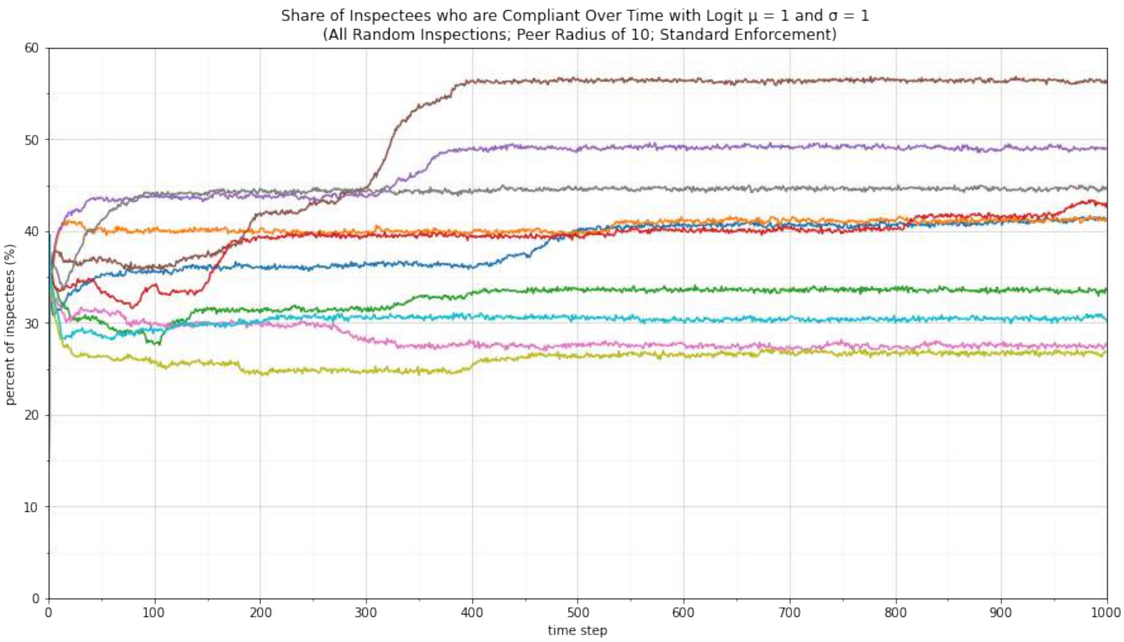

(j) $\sigma = 1$

Figure A.49: Share of compliant inspectees with varying σ & μ=1 (SE, AR, r_{peers}=10, continued).

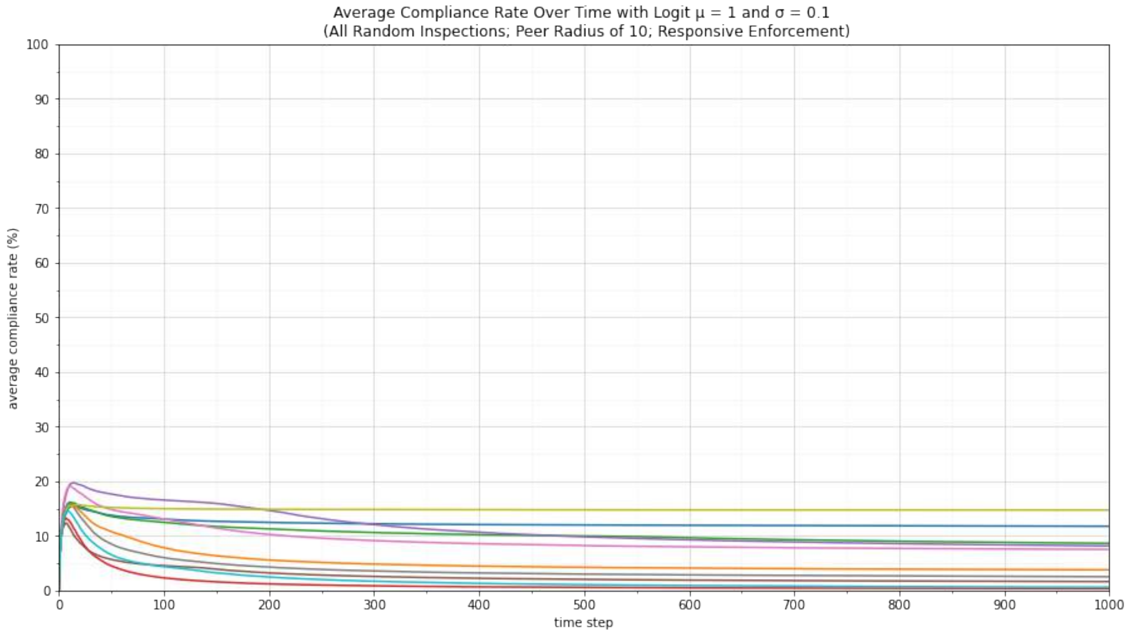

(a) $\sigma = 0.1$

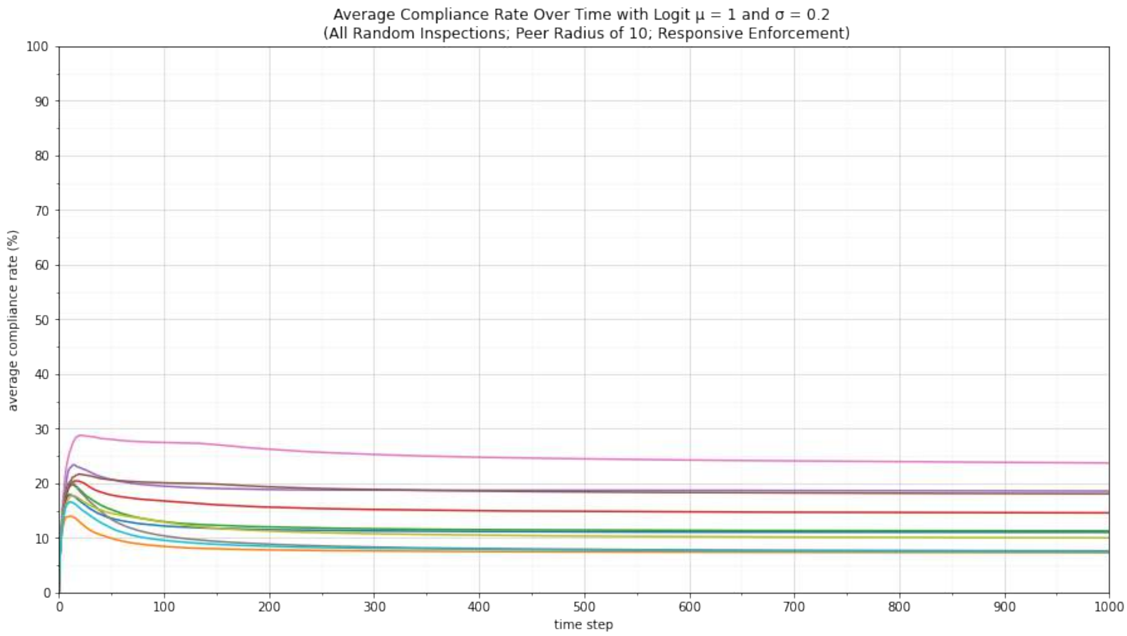

(b) $\sigma = 0.2$

Figure A.50: Average compliance rate with varying σ & μ=1 (RE, AR, r_{peers}=10).

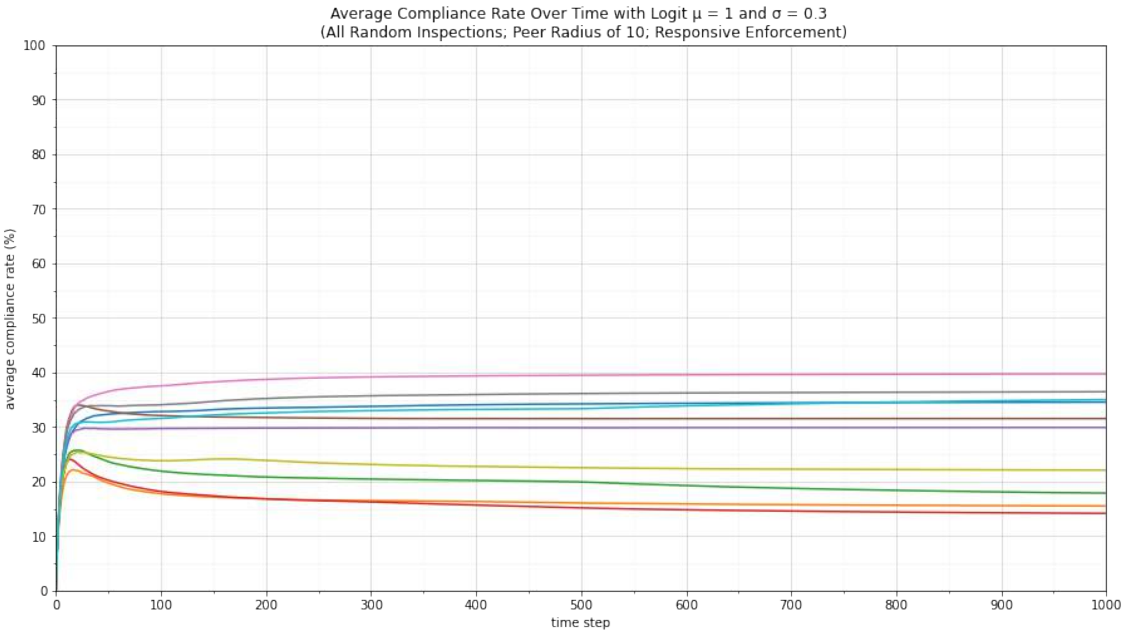

(c) $\sigma = 0.3$

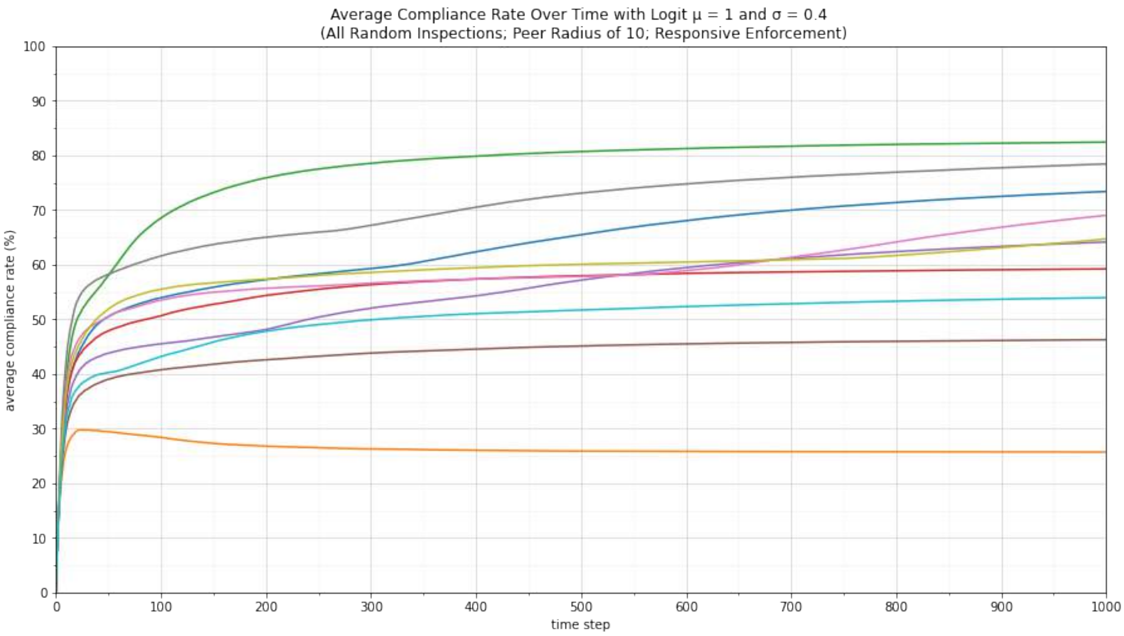

(d) $\sigma = 0.4$

Figure A.50: Average compliance rate with varying σ & μ=1 (RE, AR, r_{peers}=10, continued).

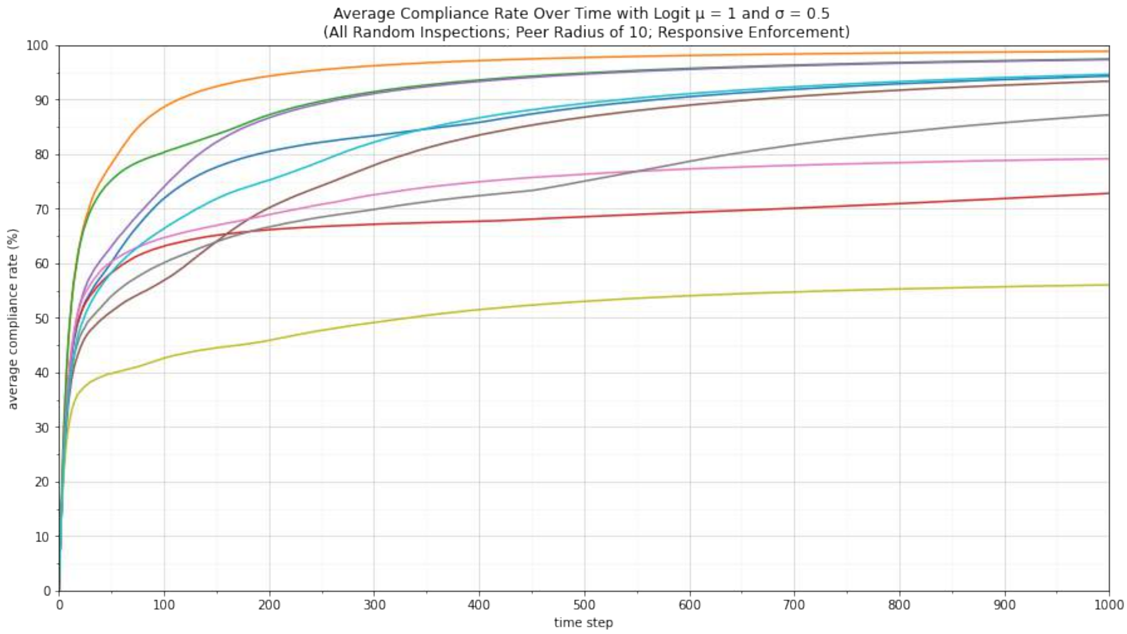

(e) $\sigma = 0.5$

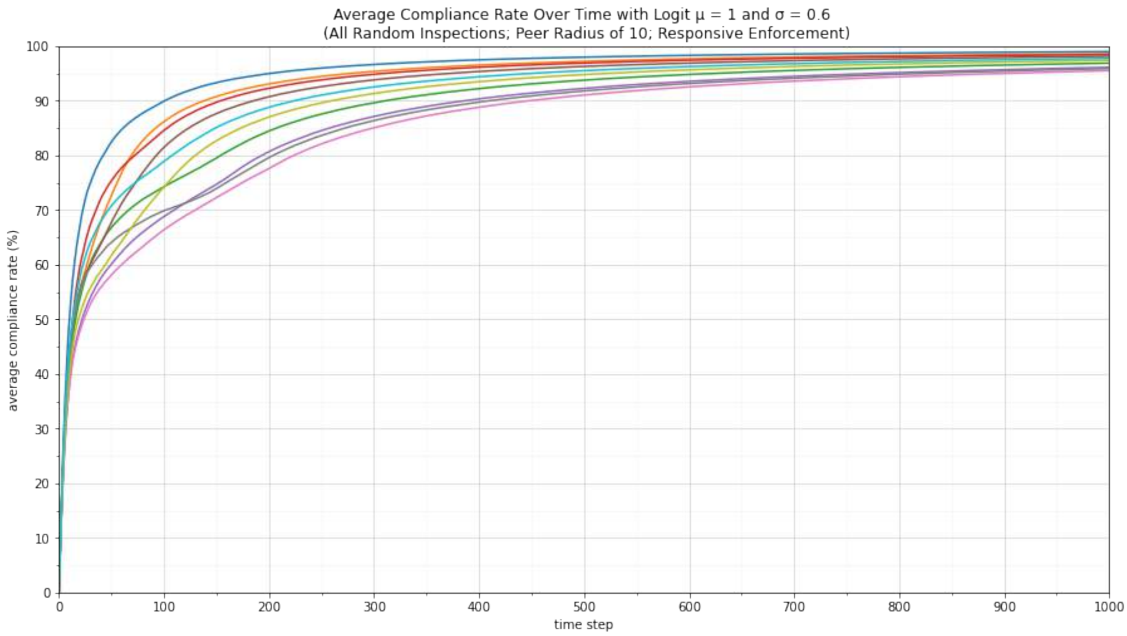

(f) $\sigma = 0.6$

Figure A.50: Average compliance rate with varying σ & μ=1 (RE, AR, r_{peers}=10, continued).

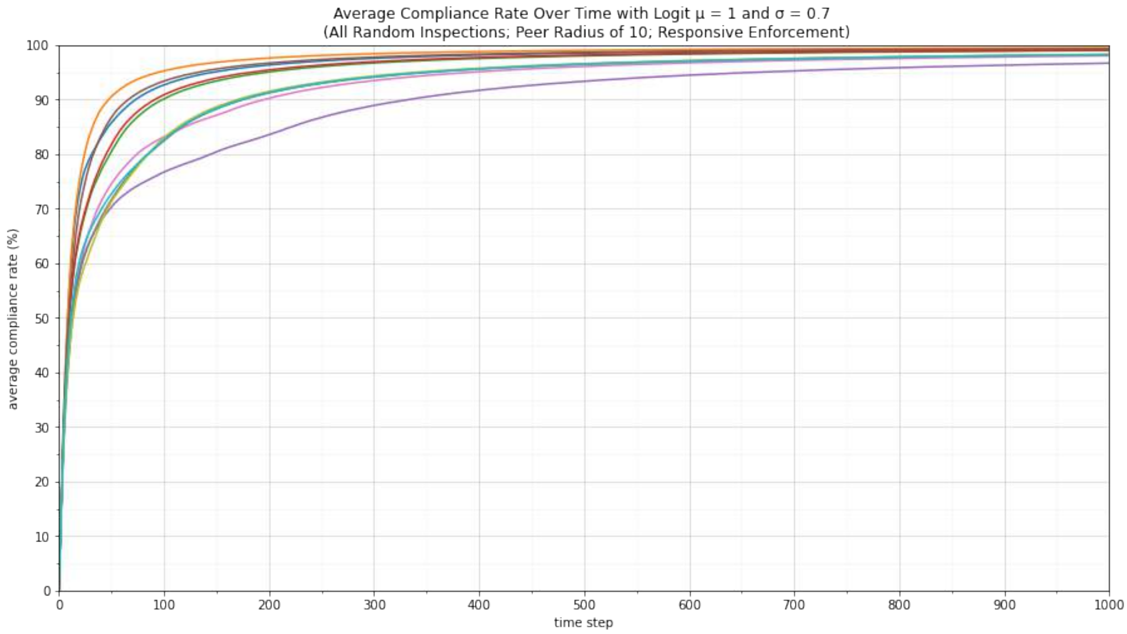

(g) σ = 0.7

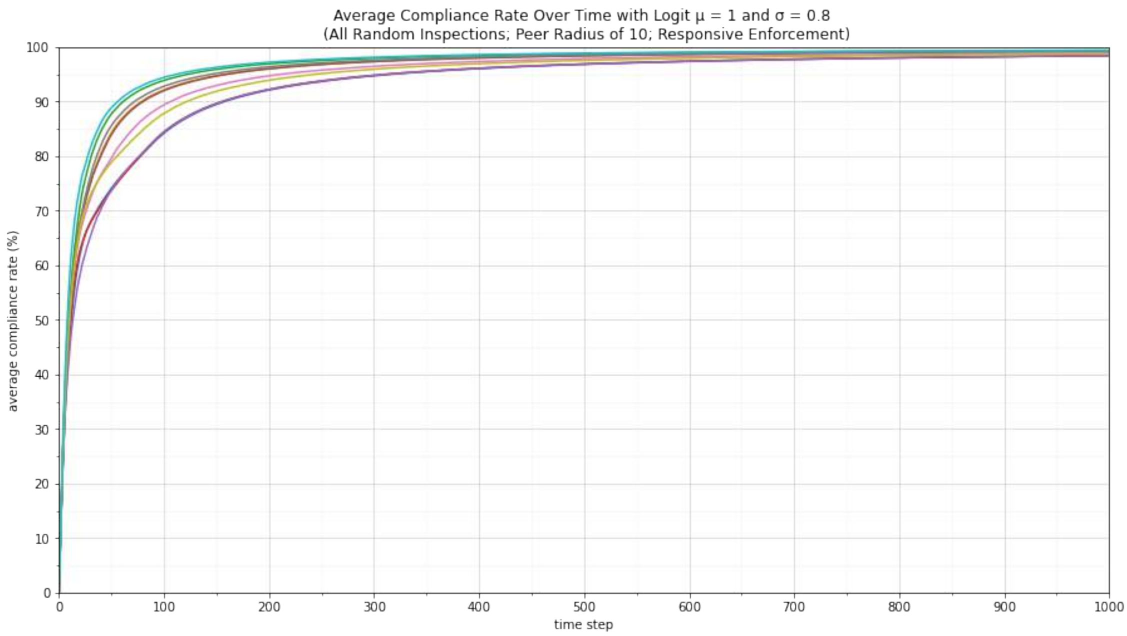

(h) σ = 0.8

Figure A.50: Average compliance rate with varying σ & μ=1 (RE, AR, r_{peers}=10, continued).

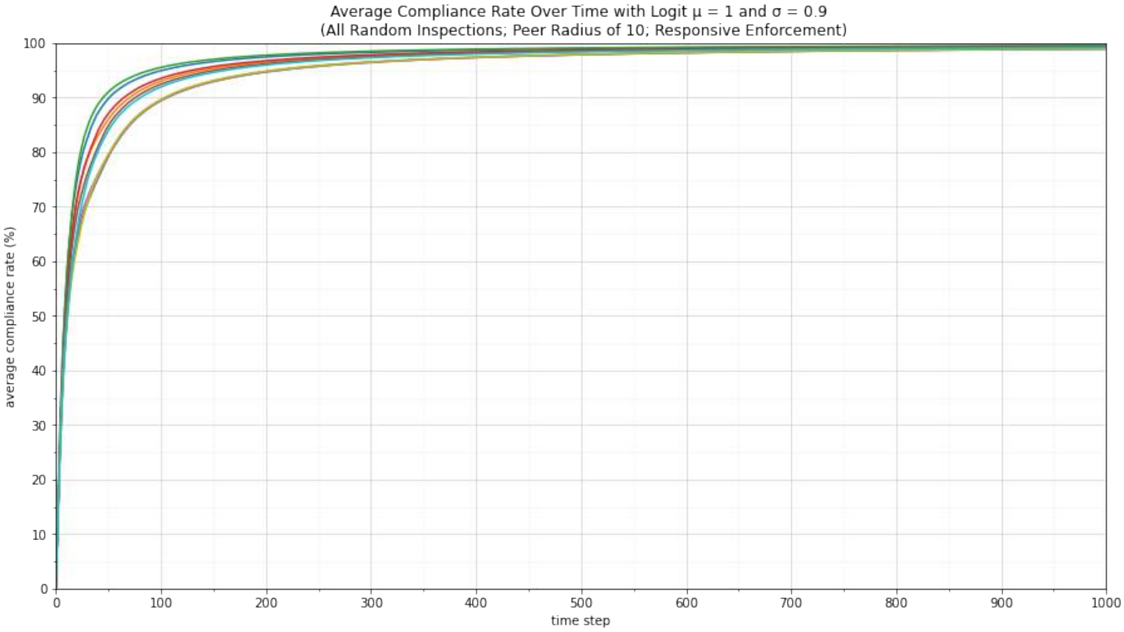

(i) $\sigma = 0.9$

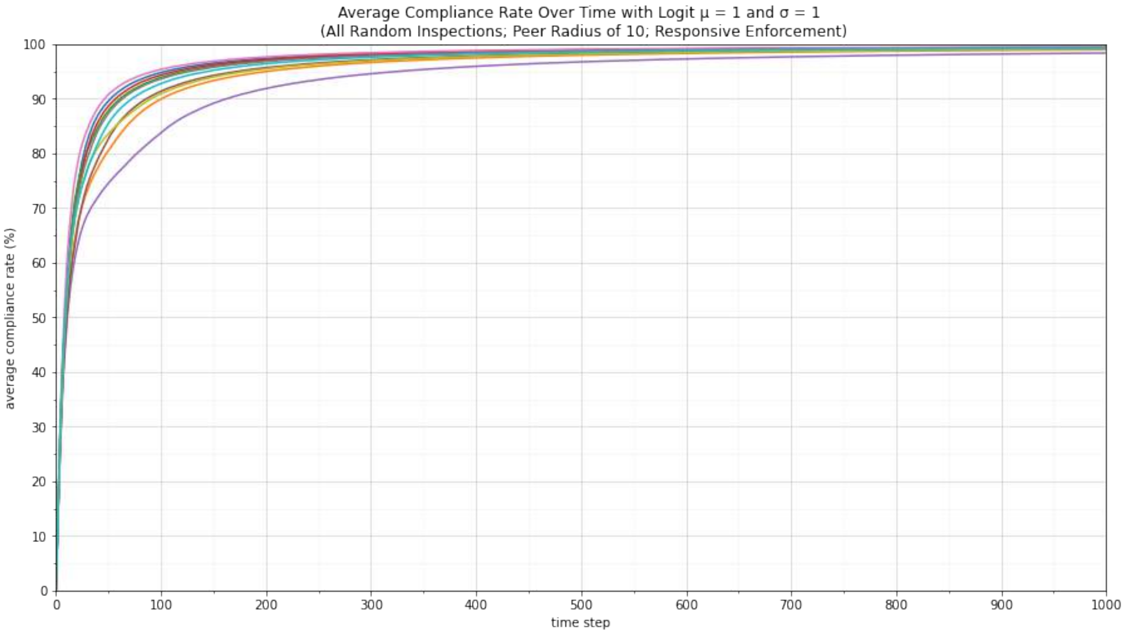

(j) $\sigma = 1$

Figure A.50: Average compliance rate with varying σ & μ=1 (RE, AR, r_{peers}=10, continued).

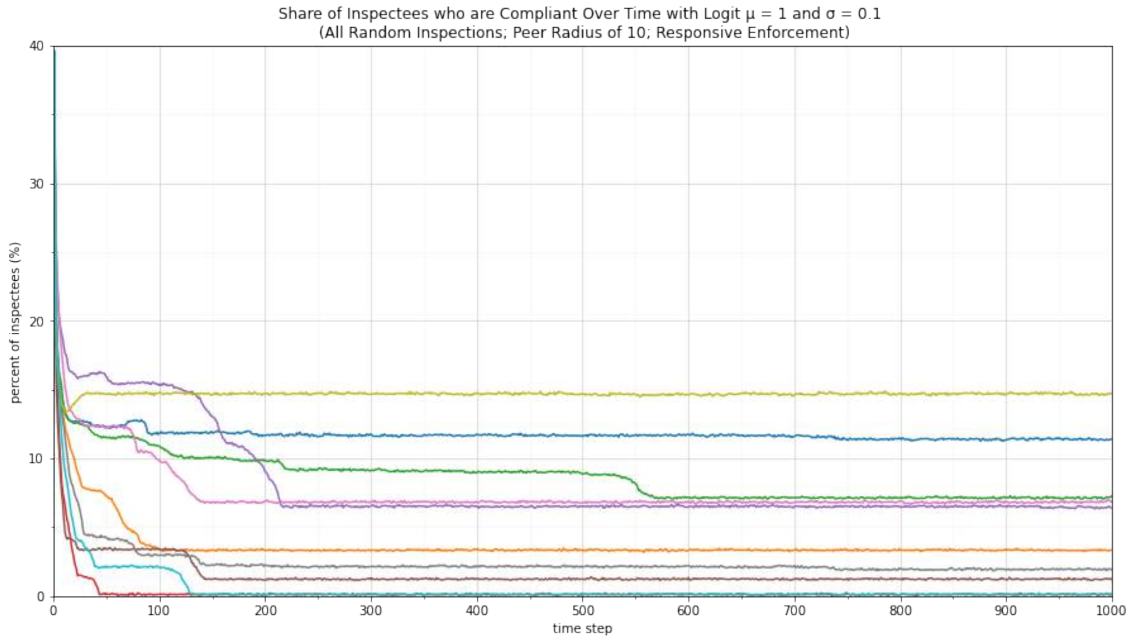

(a) σ = 0.1

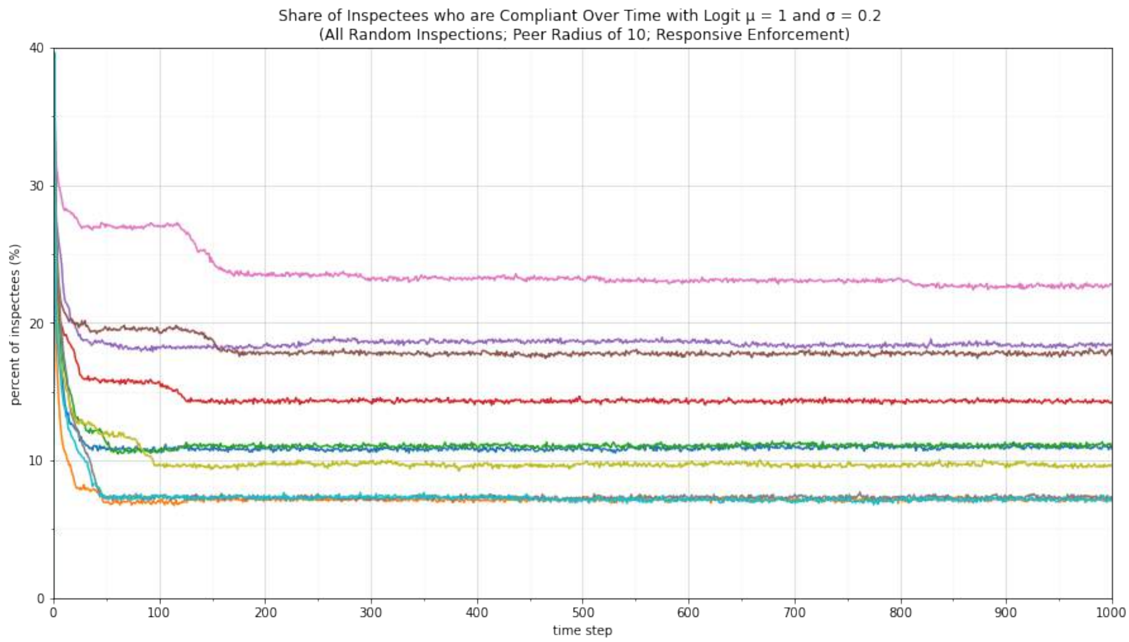

(b) σ = 0.2

Figure A.51: Share of compliant inspectees with varying σ & μ=1 (RE, AR, r_{peers}=10).

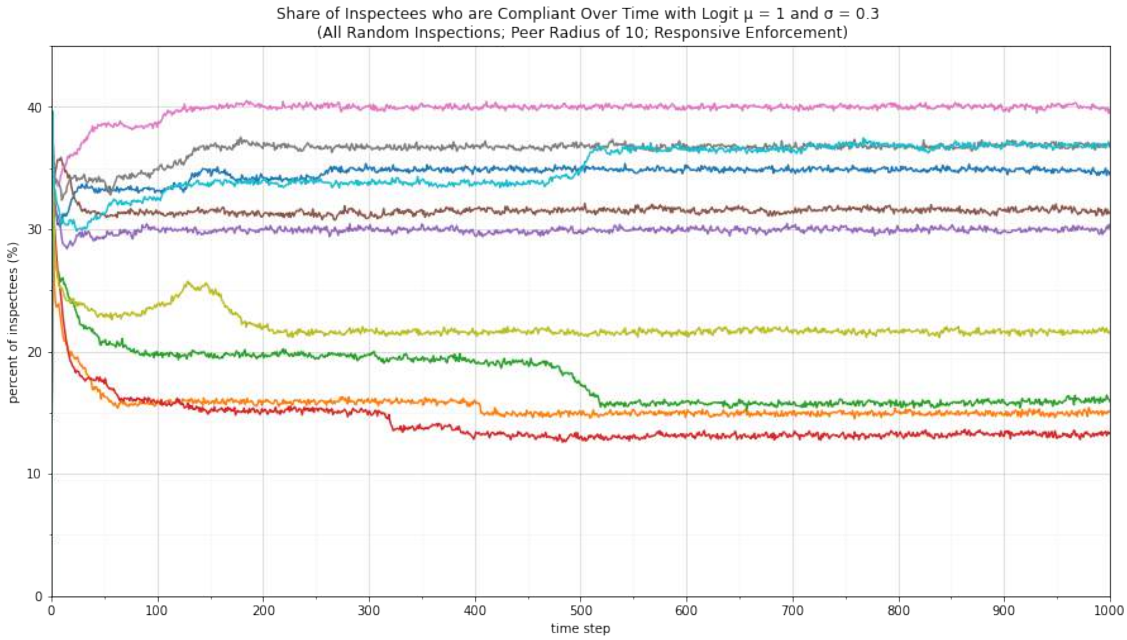

(c) $\sigma = 0.3$

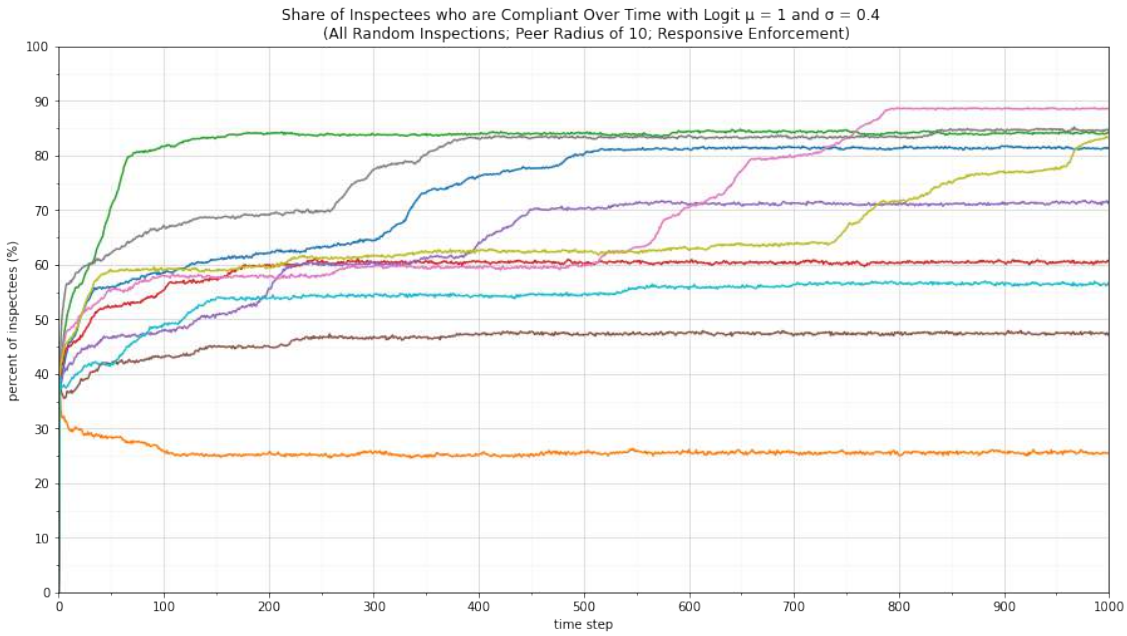

(d) $\sigma = 0.4$

Figure A.51: Share of compliant inspectees with varying σ & μ=1 (RE, AR, r_{peers}=10, continued).

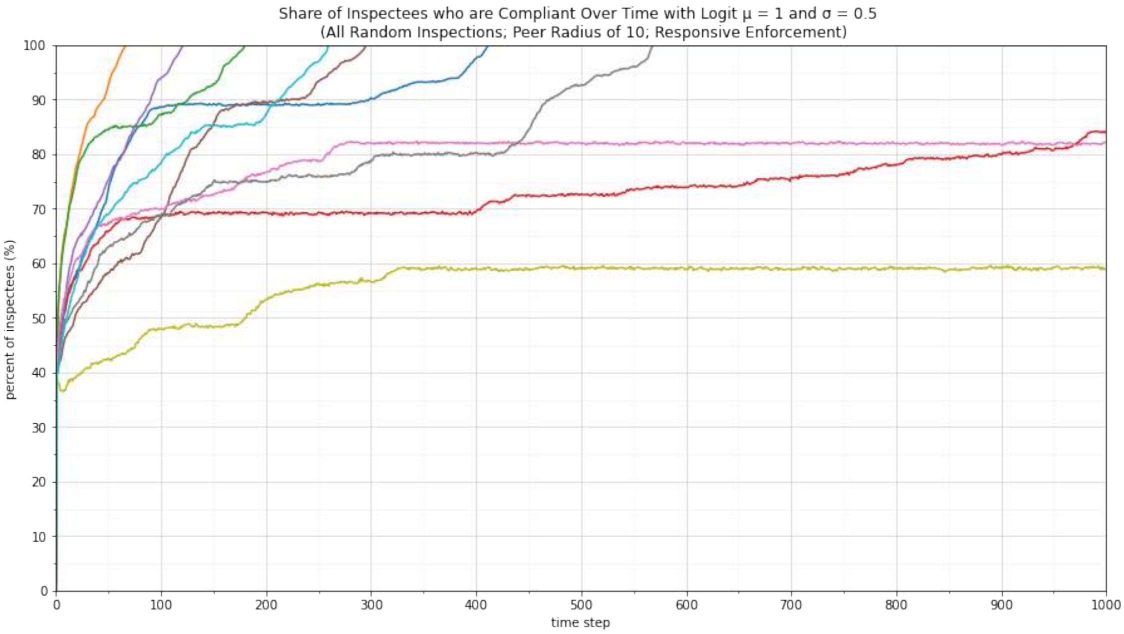

(e) $\sigma = 0.5$

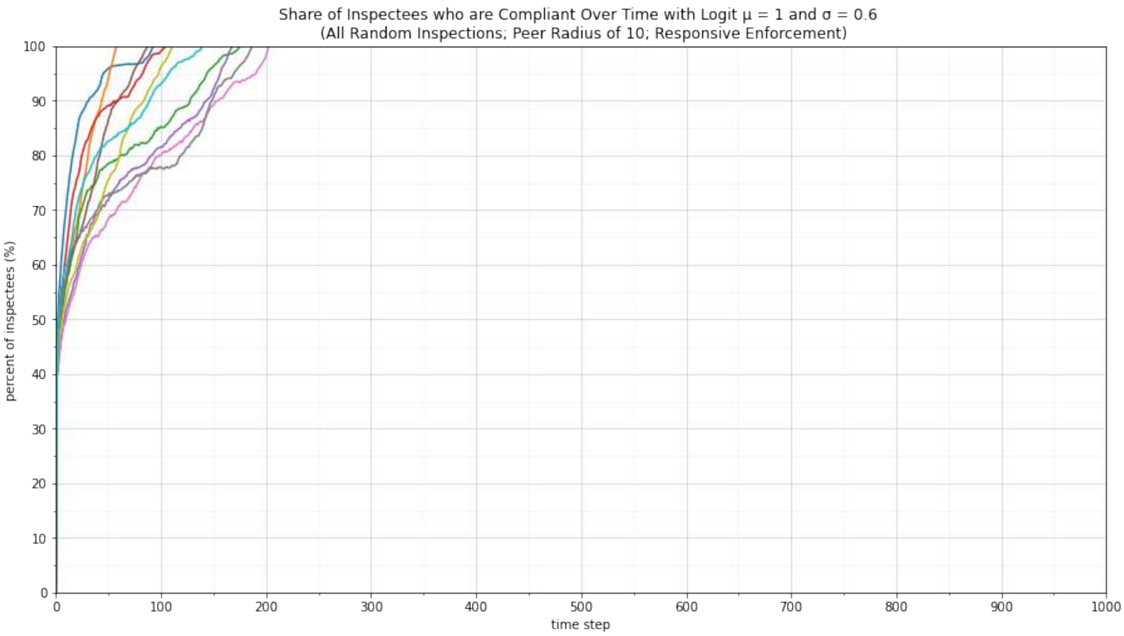

(f) $\sigma = 0.6$

Figure A.51: Share of compliant inspectees with varying σ & μ=1 (RE, AR, r_{peers}=10, continued).

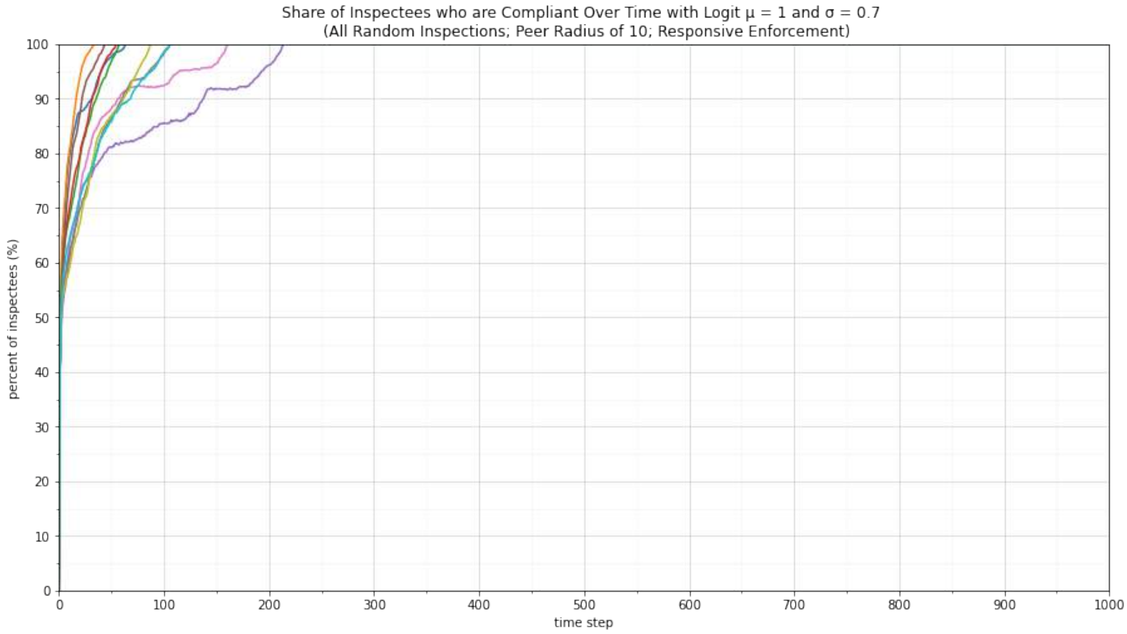

(g) $\sigma = 0.7$

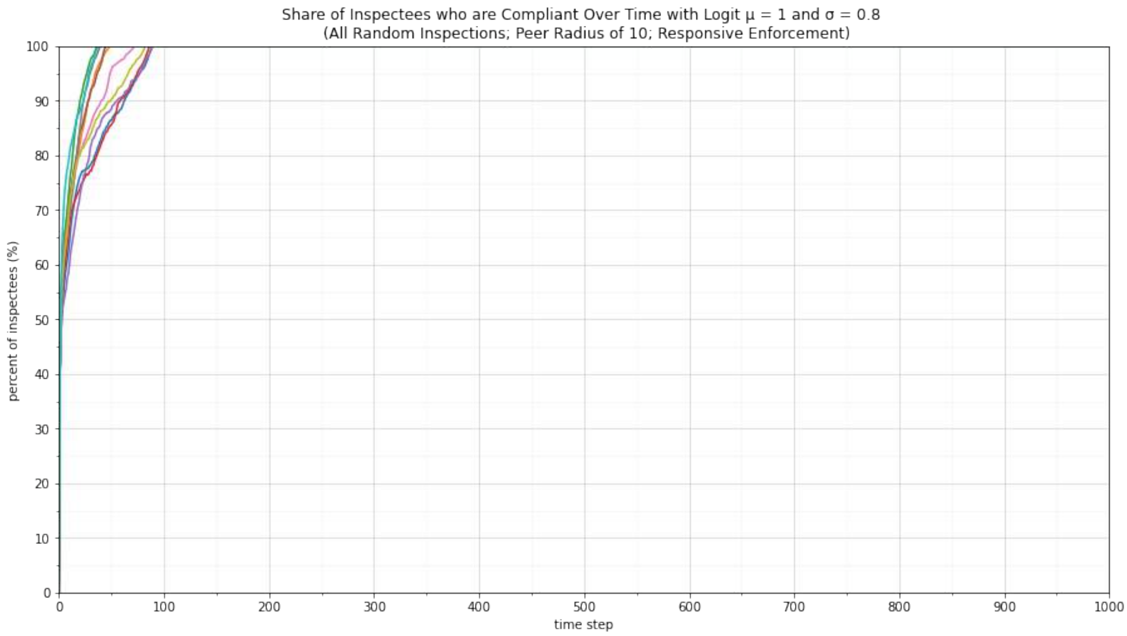

(h) $\sigma = 0.8$

Figure A.51: Share of compliant inspectees with varying σ & μ=1 (RE, AR, r_{peers}=10, continued).

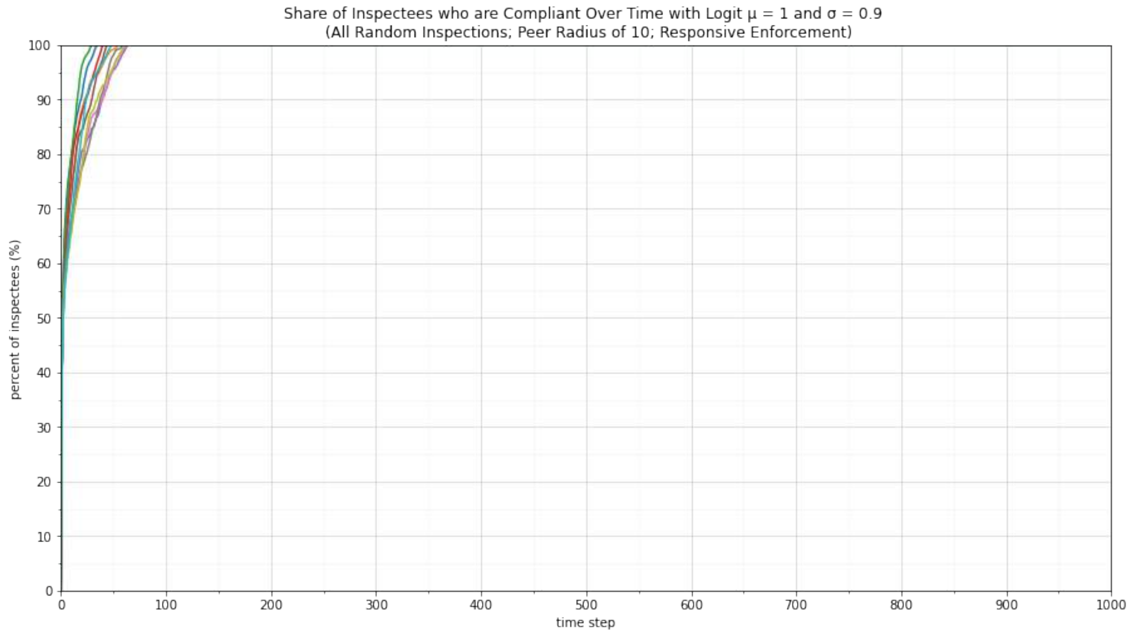

(i) $\sigma = 0.9$

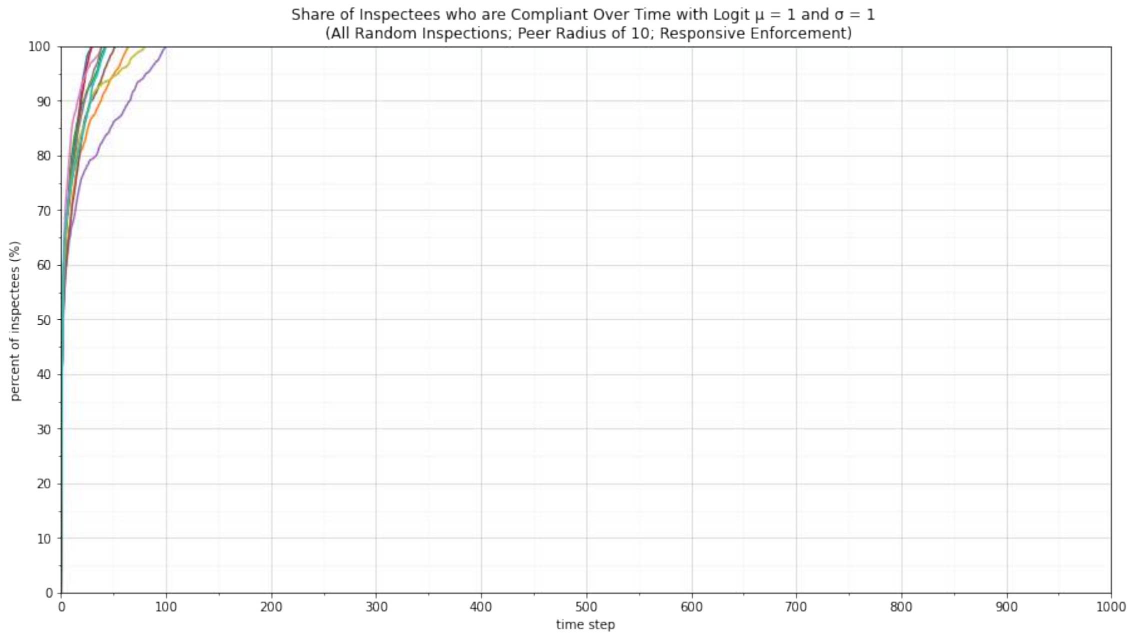

(j) $\sigma = 1$

Figure A.51: Share of compliant inspectees with varying σ & μ=1 (RE, AR, r_{peers}=10, continued).

Varying Logit μ

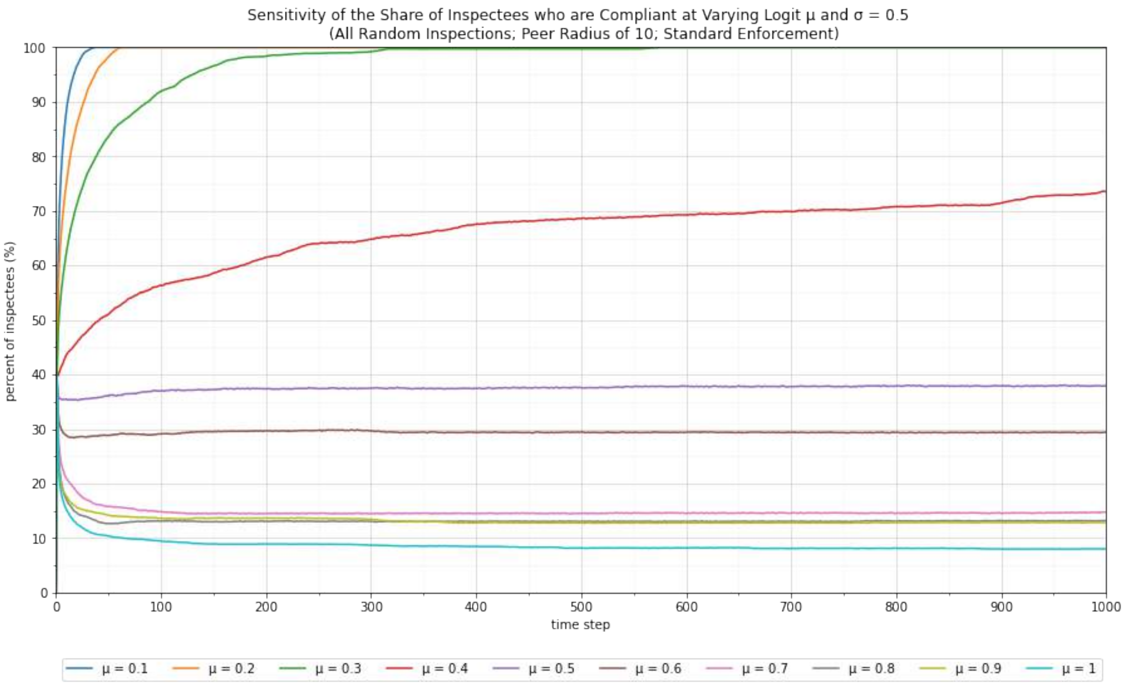

(a) Standard enforcement

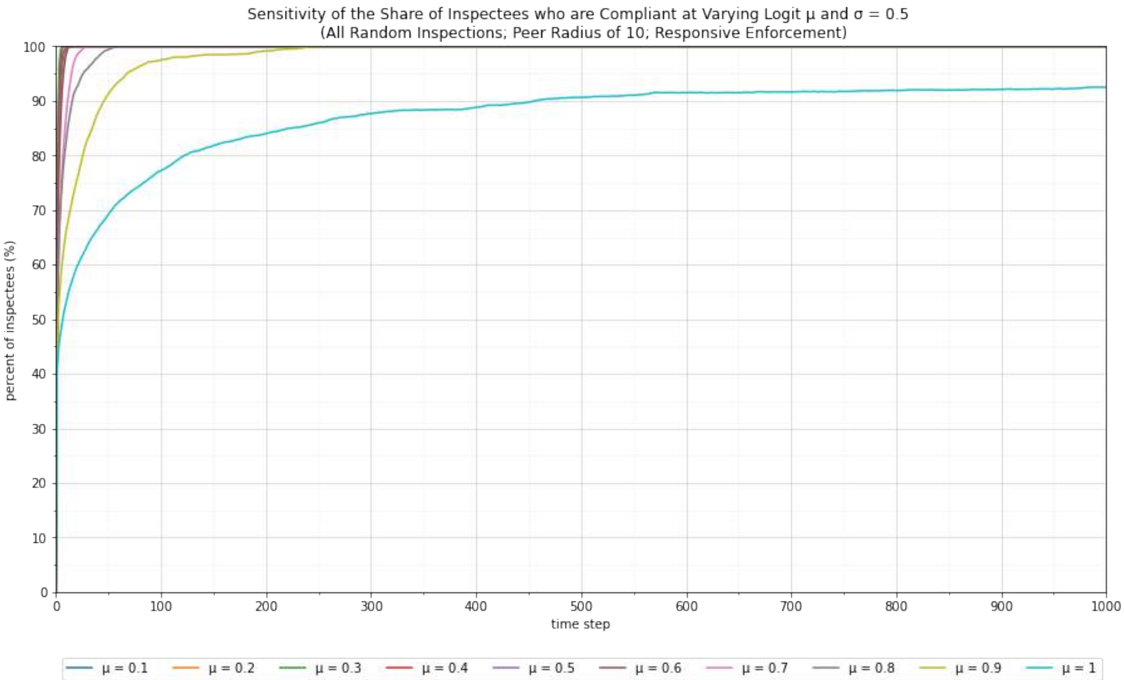

(b) Responsive enforcement.

Figure A.52: Share of compliant inspectees with varying σ & μ=1 (AR, r_{peers}=10).

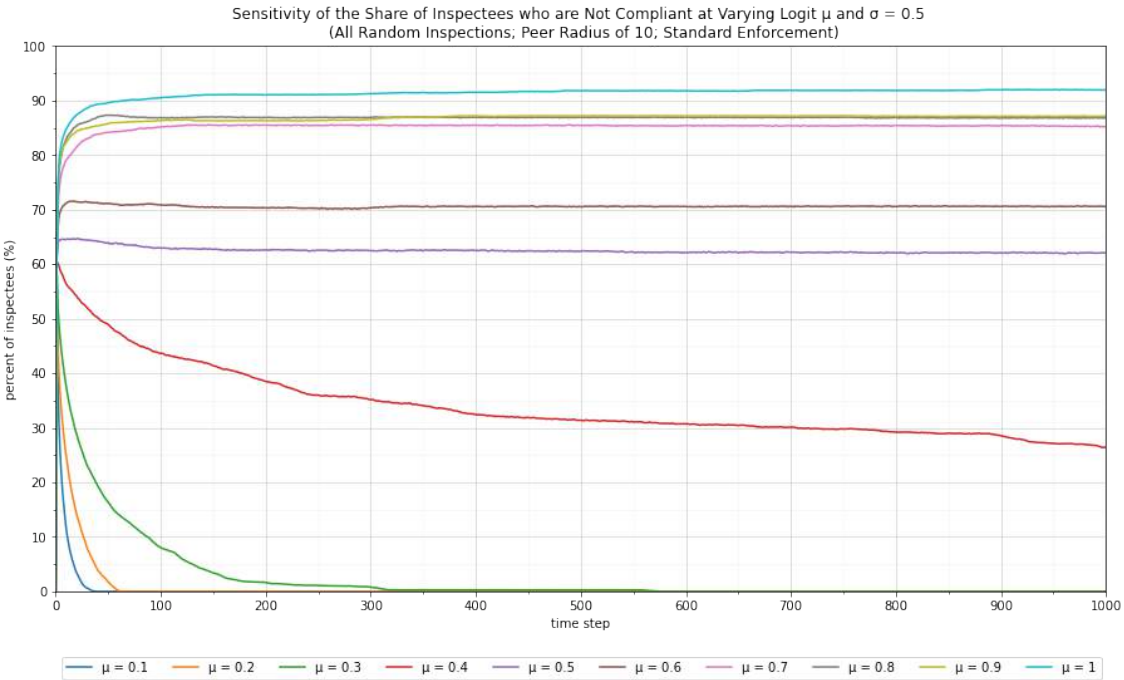

(a) Standard enforcement.
For a detailed breakdown by compliance level, see Figure A.54.

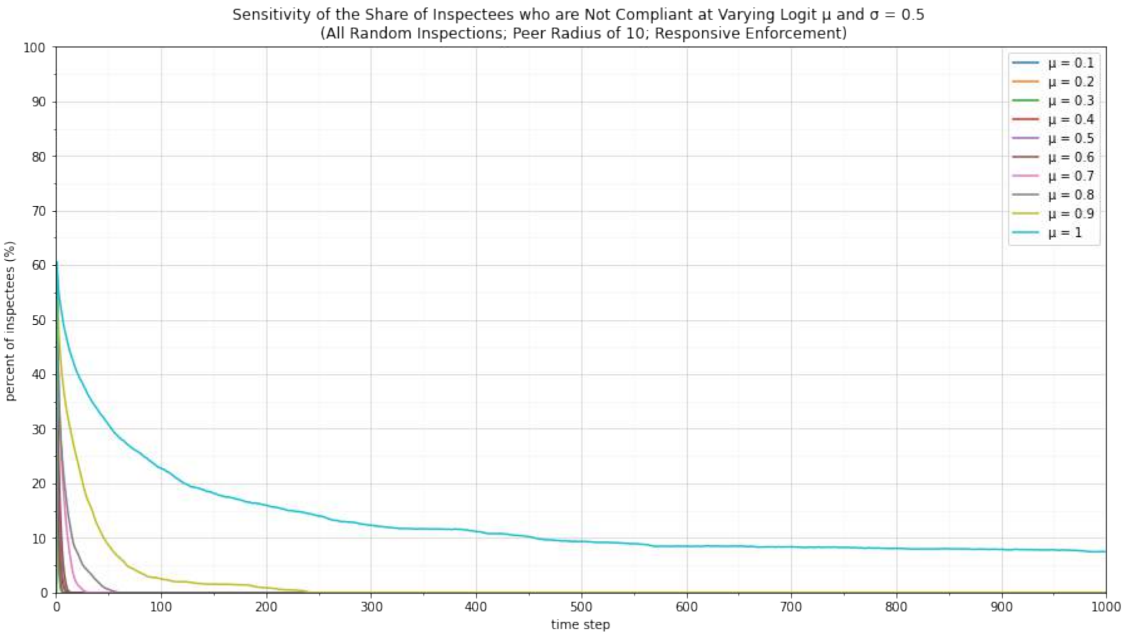

(b) Responsive enforcement.
For a detailed breakdown by compliance level, see Figure A.55.

Figure A.53: Share of non-compliant inspectees with varying σ & μ=1 (AR, r_{peers}=10).

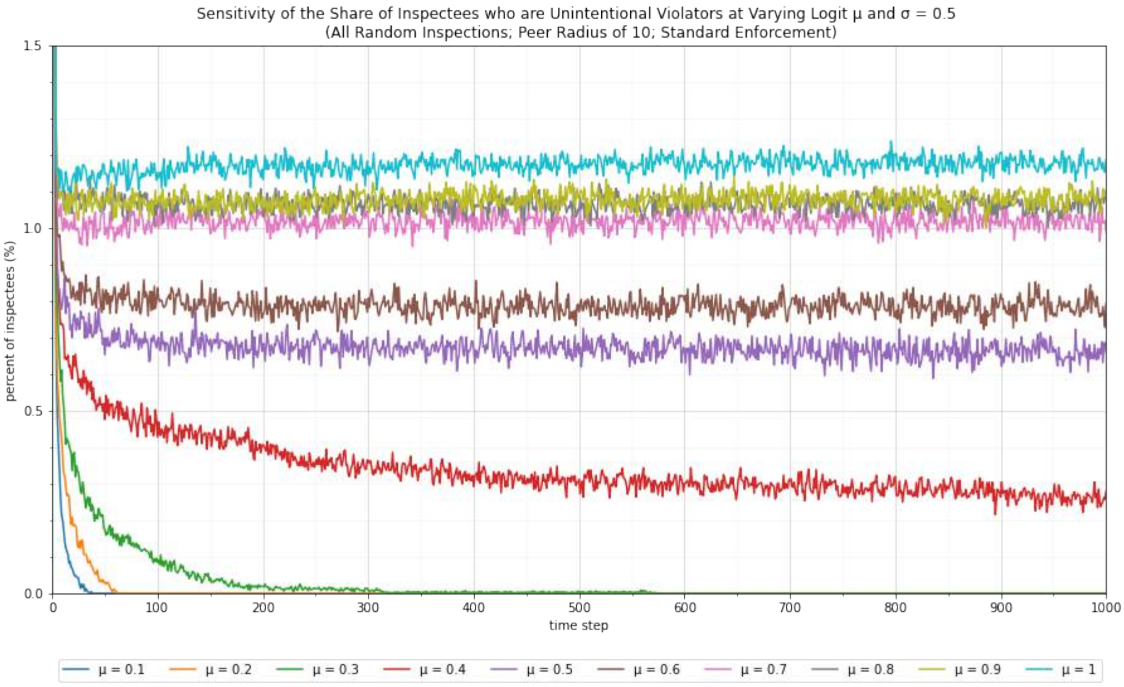

(a) Unintentional violators.

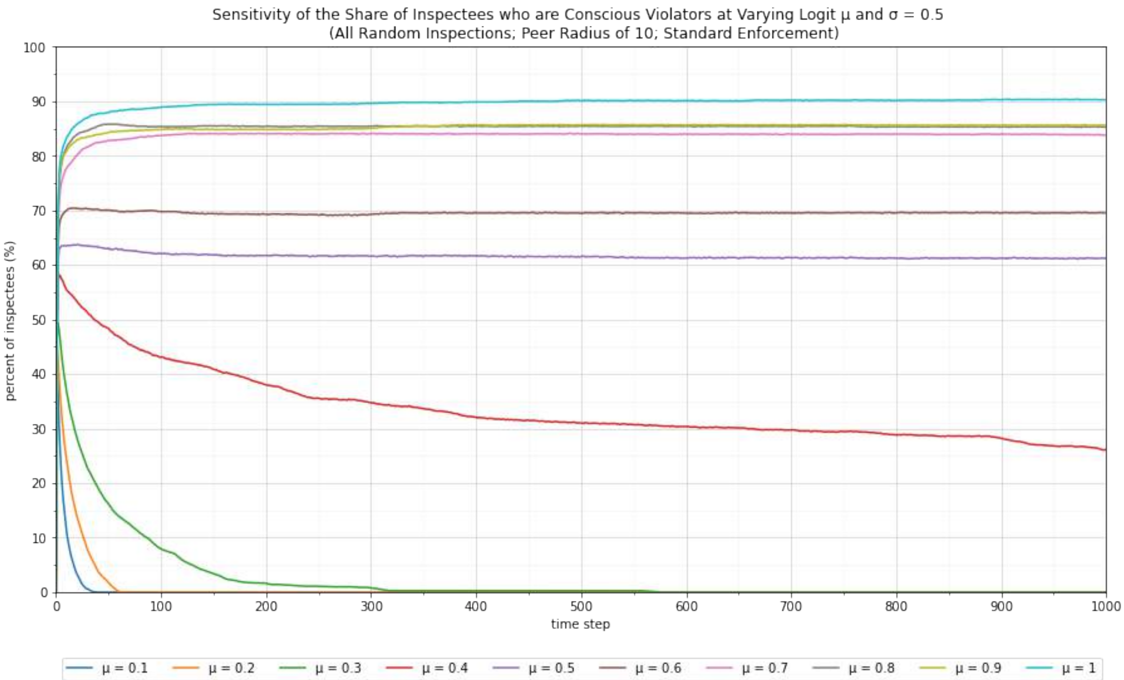

(b) Conscious violators.

Figure A.54: Non-compliant inspectee breakdown with varying σ & μ=1 (SE, AR, r_{peers}=10).

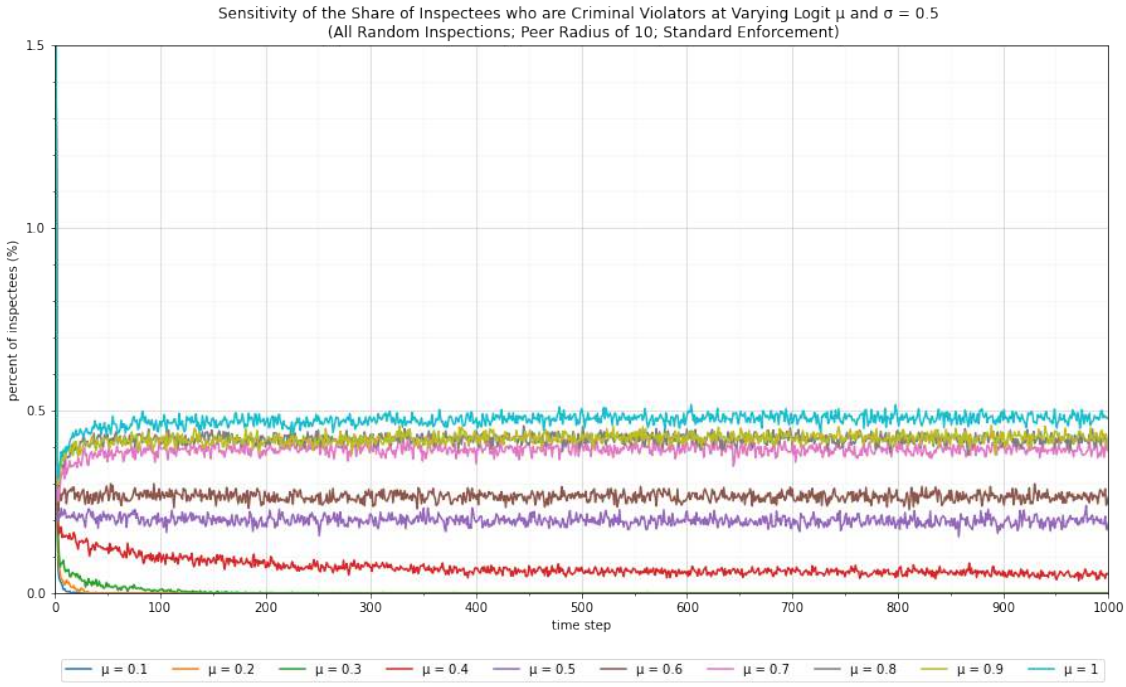

(c) Criminal violators.

Figure A.54: Non-compliant inspectee breakdown with varying σ & μ=1 (SE, AR, r_{peers}=10) (continued).

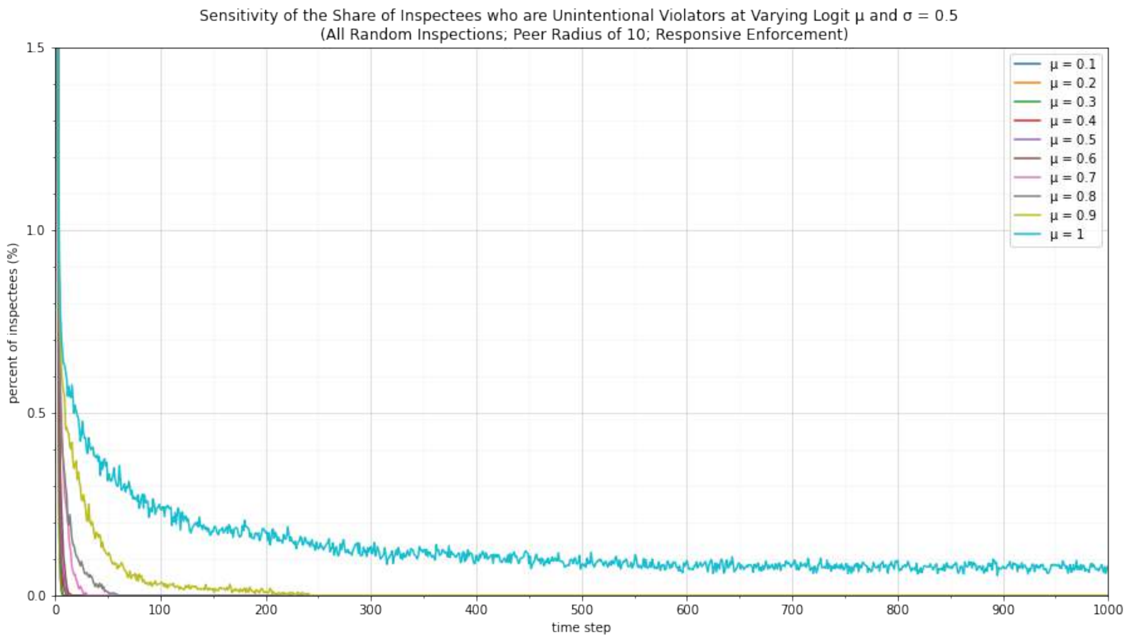

(a) Unintentional violators.

Figure A.55: Non-compliant inspectee breakdown with varying σ & μ=1 (RE, AR, r_{peers}=10).

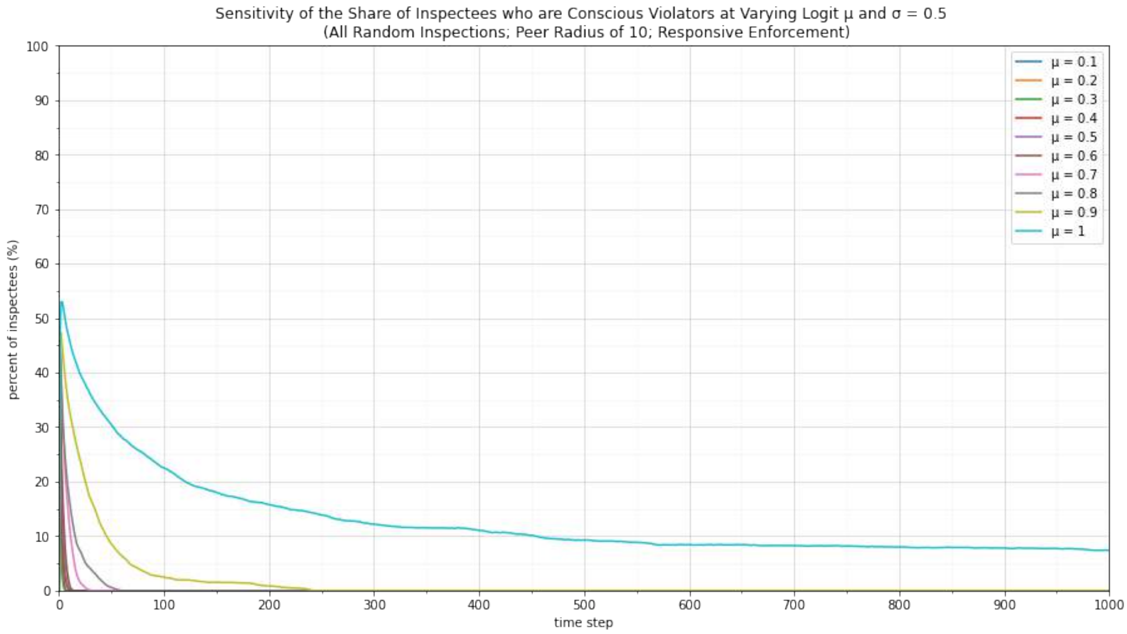

(b) Conscious violators.

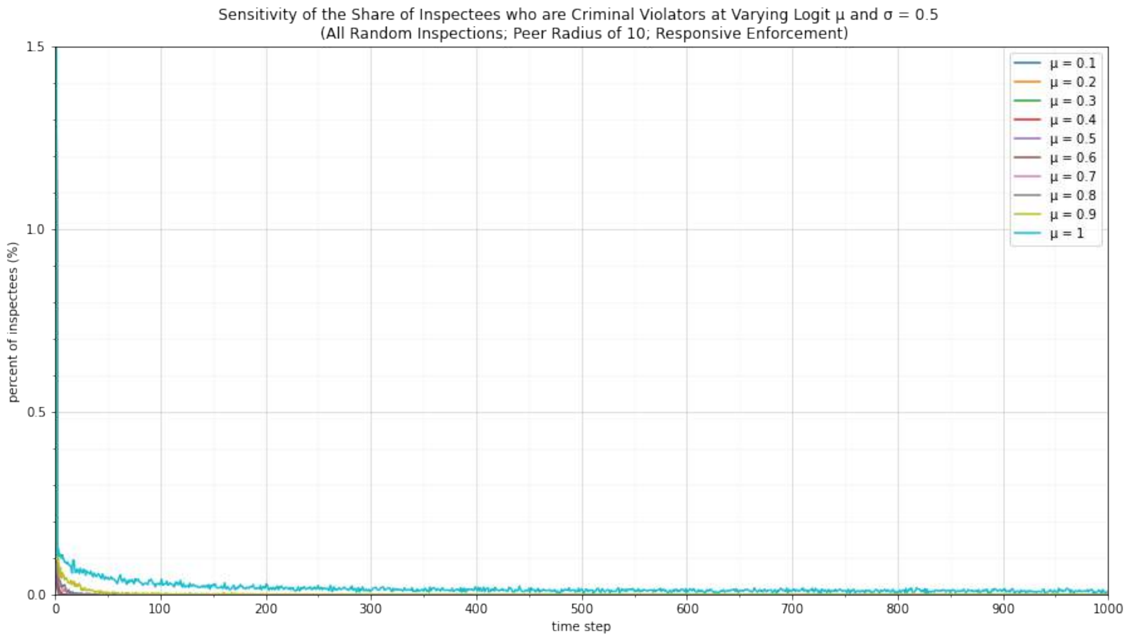

(c) Criminal violators.

Figure A.55: Non-compliant inspectee breakdown with varying σ & μ=1 (RE, AR, r_{peers}=10, continued).

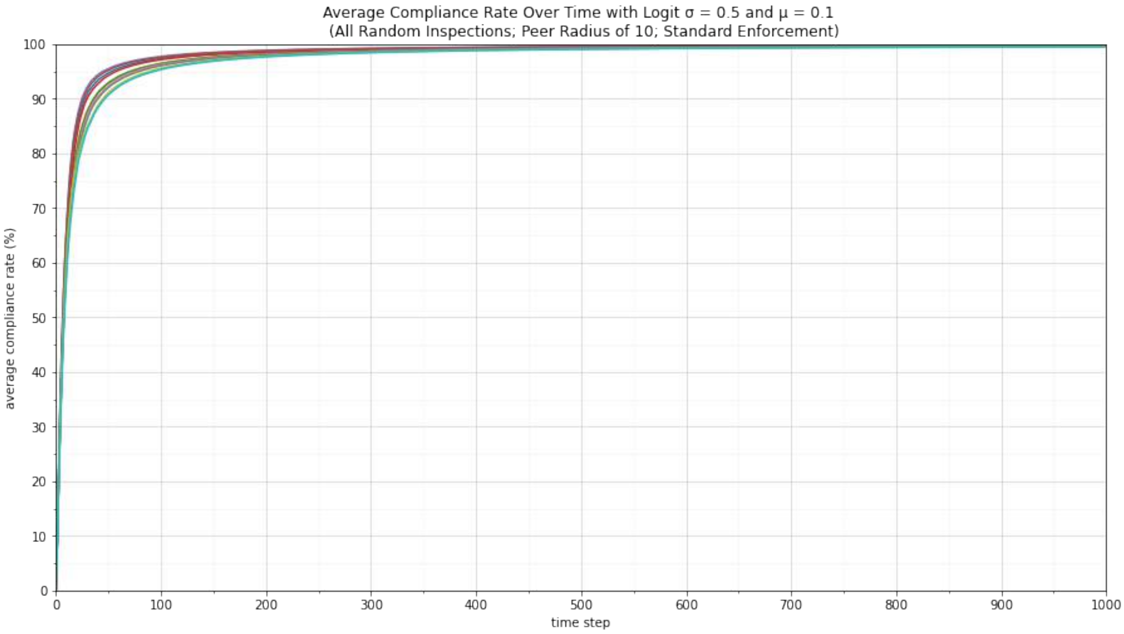

(a) $\sigma = 0.1$

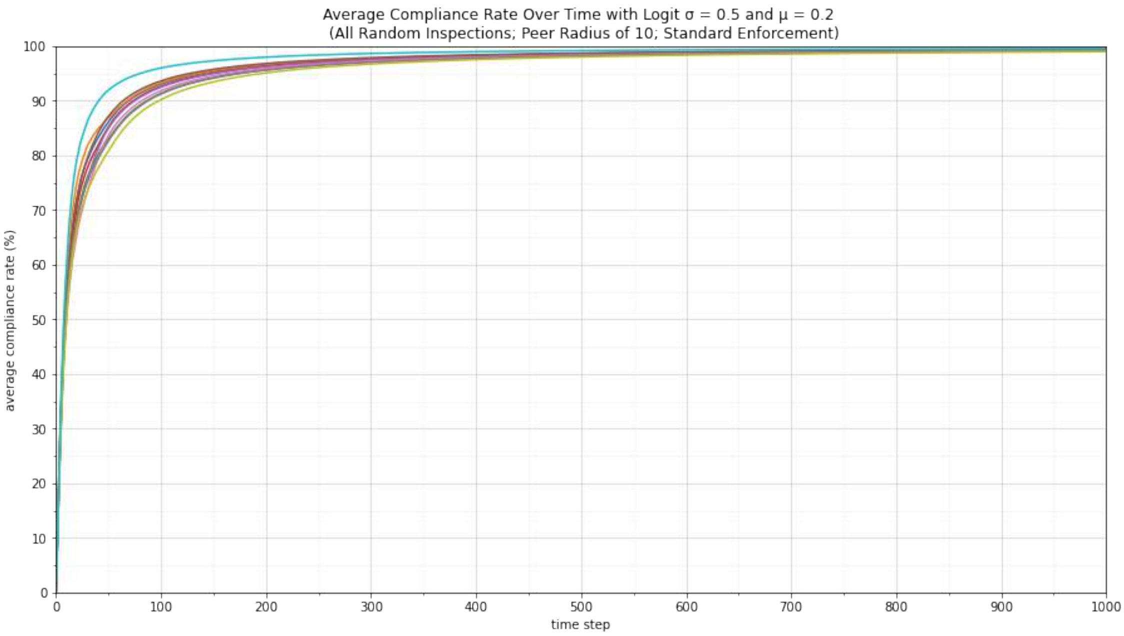

(b) $\sigma = 0.2$

Figure A.56: Average compliance rate with varying σ & μ=1 (SE, AR, r_{peers}=10).

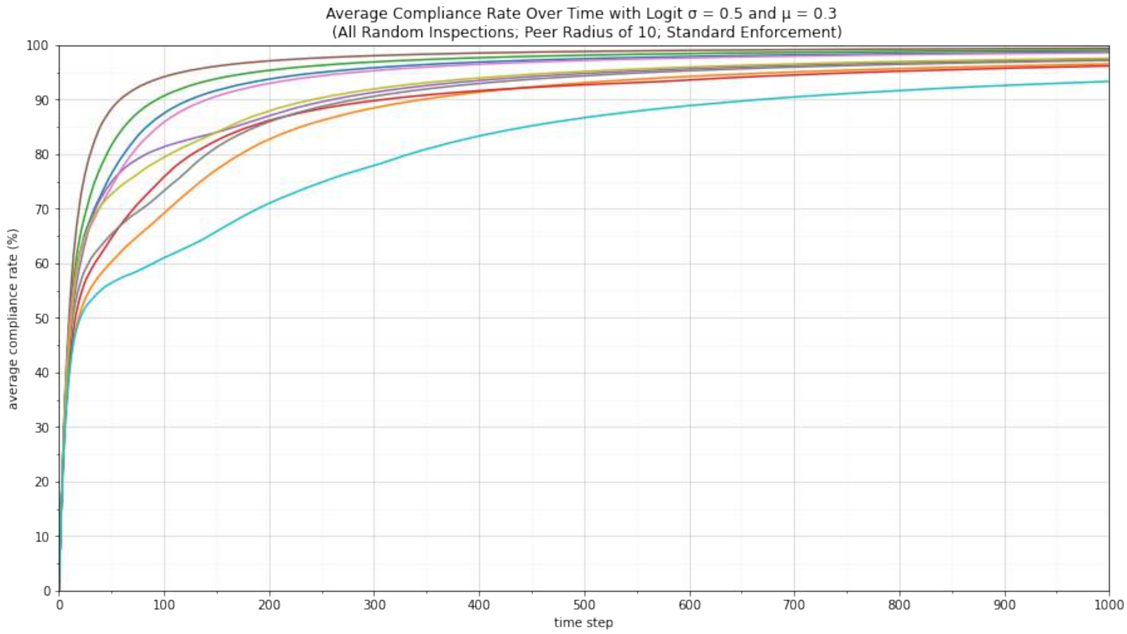

(c) σ = 0.3

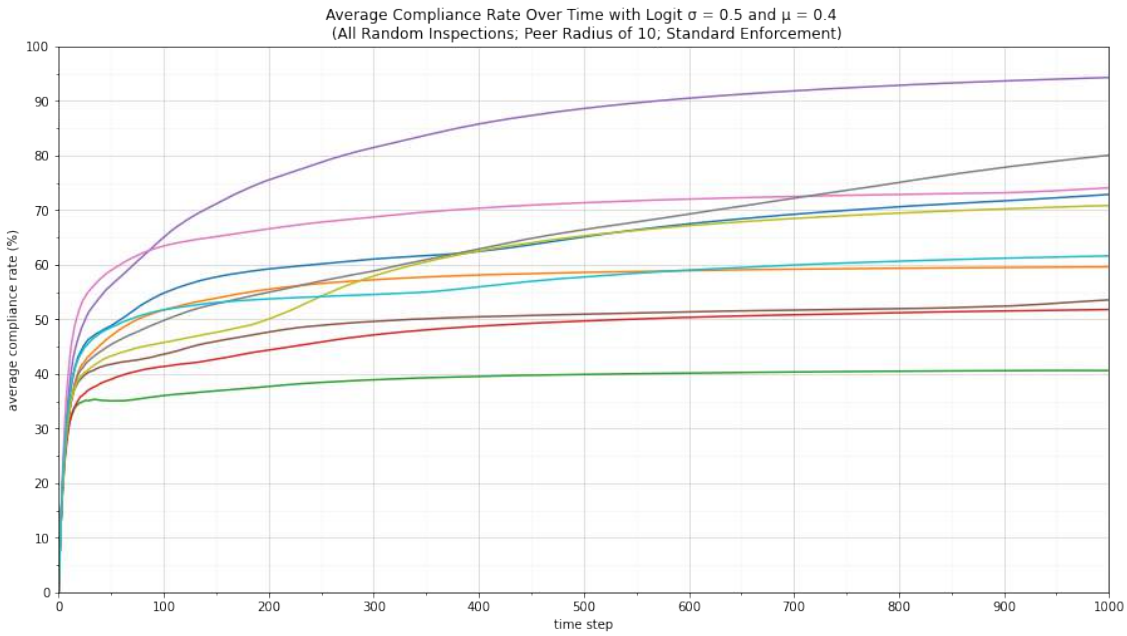

(d) σ = 0.4

Figure A.56: Average compliance rate with varying σ & μ=1 (SE, AR, r_{peers}=10, continued).

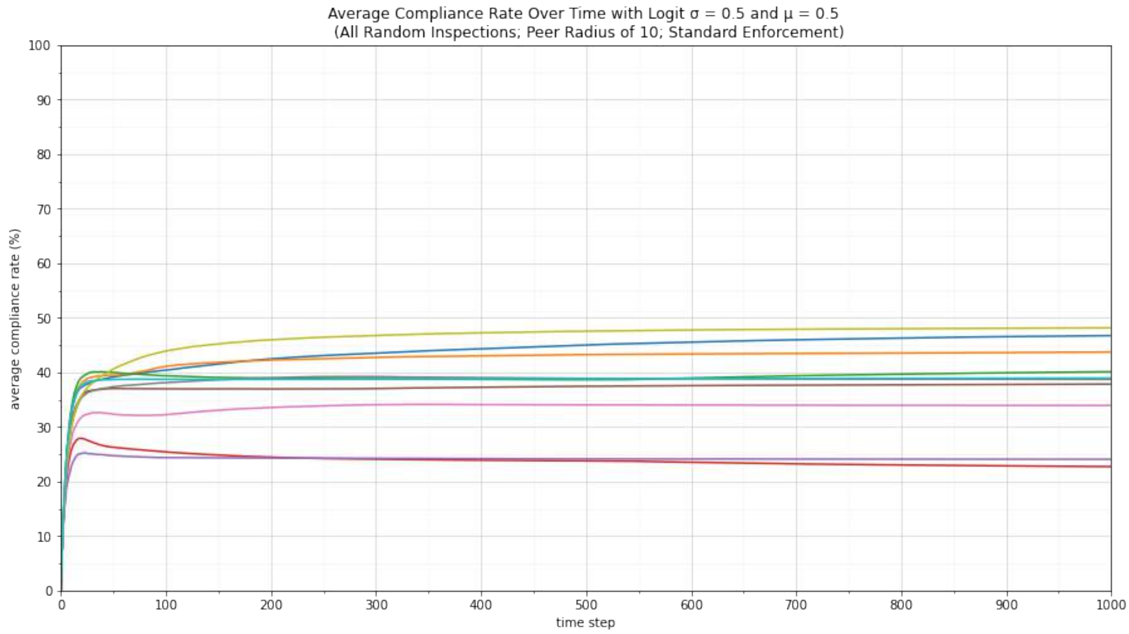

(e) $\sigma = 0.5$

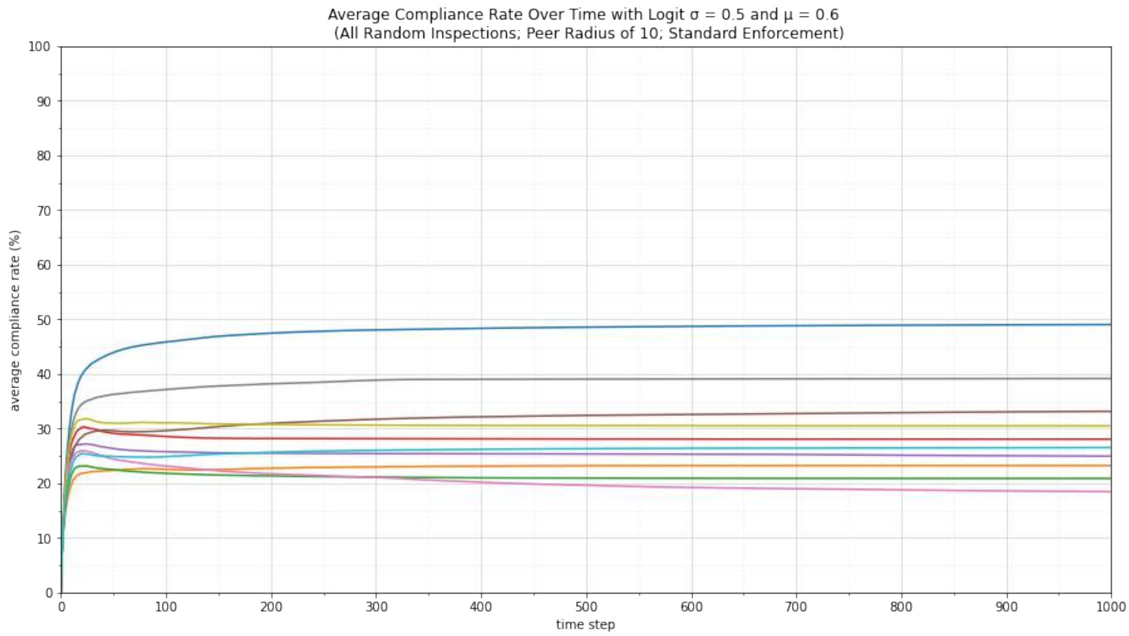

(f) $\sigma = 0.6$

Figure A.56: Average compliance rate with varying σ & μ=1 (SE, AR, r_{peers}=10, continued).

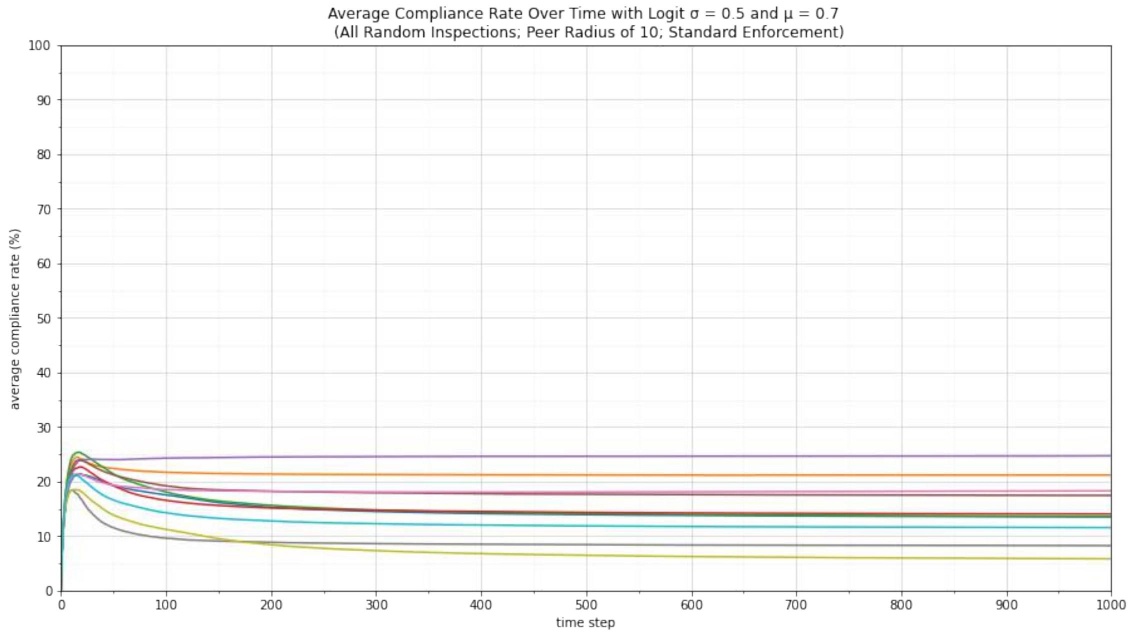

(g) σ = 0.7

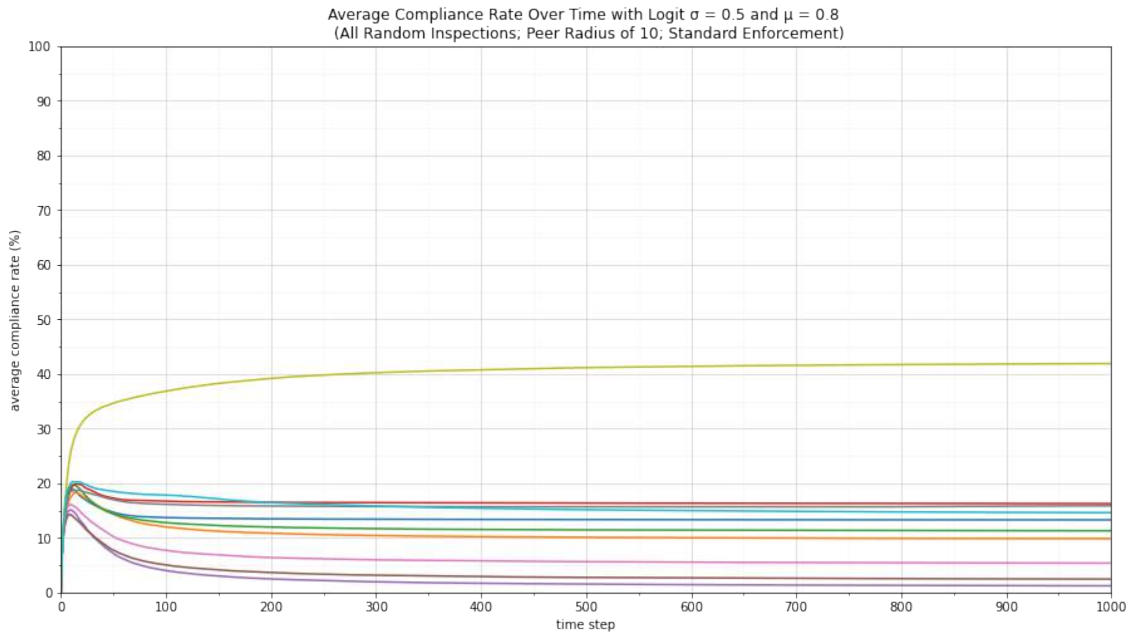

(h) σ = 0.8

Figure A.56: Average compliance rate with varying σ & μ=1 (SE, AR, r_{peers}=10, continued).

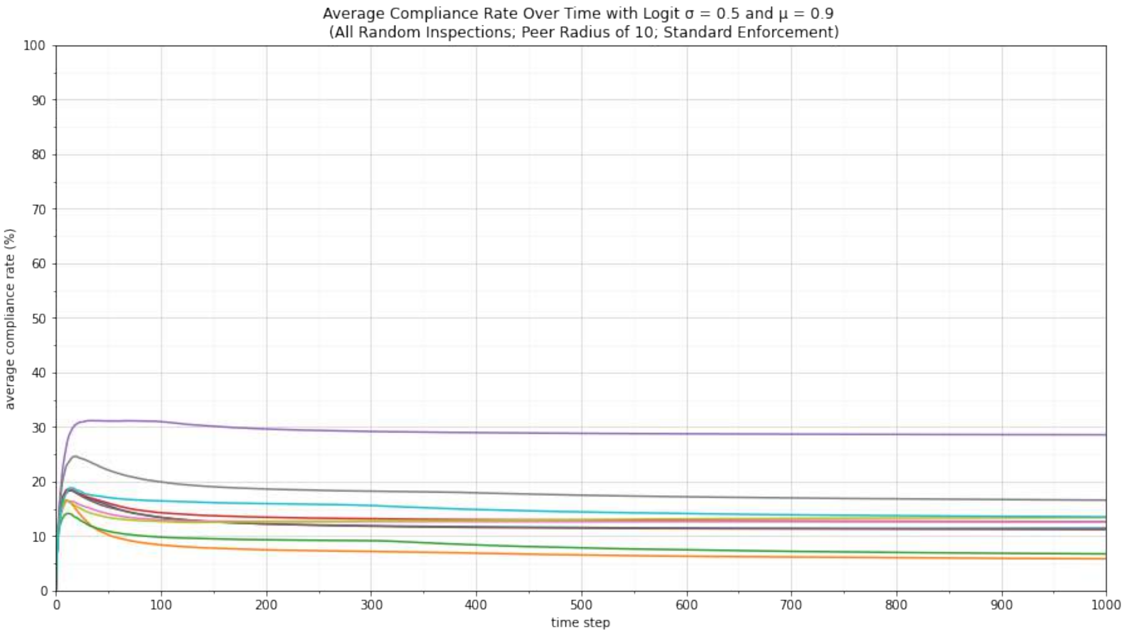

(i) σ = 0.9

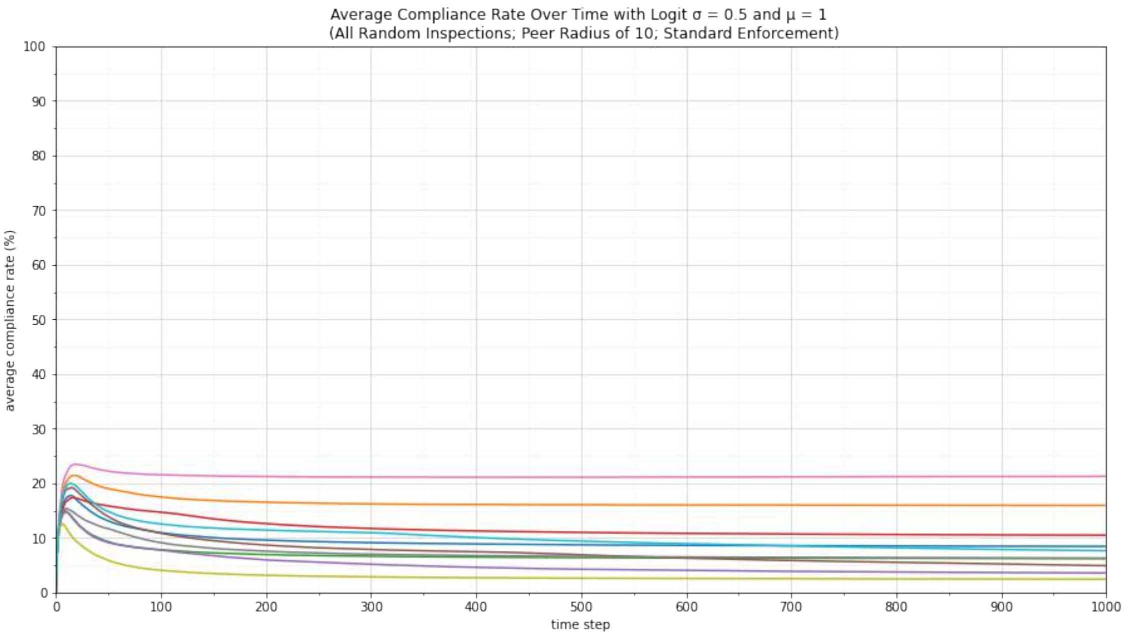

(j) σ = 1

Figure A.56: Average compliance rate with varying σ & μ=1 (SE, AR, r_{peers}=10, continued).

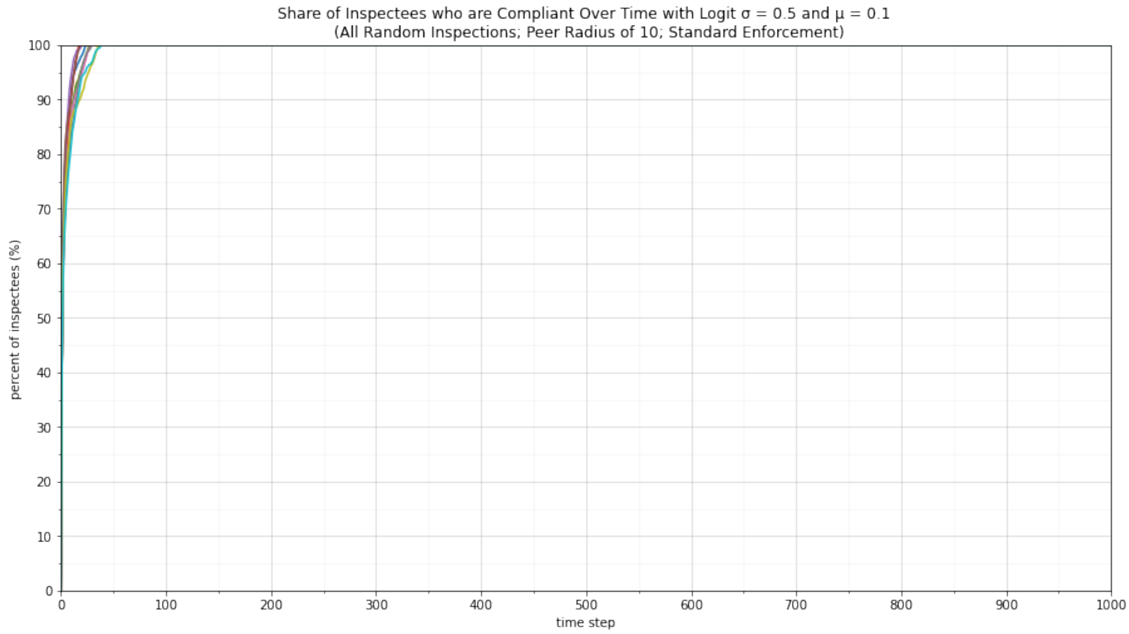

(a) σ = 0.1

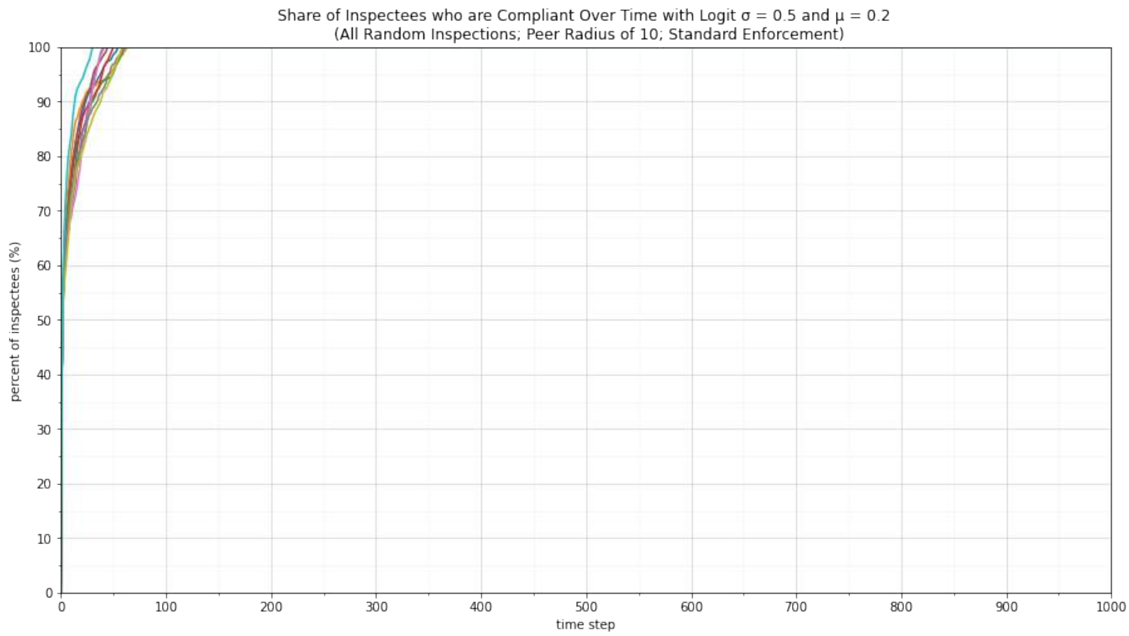

(b) σ = 0.2

Figure A.57: Average compliance rate with varying σ & μ=1 (SE, AR, r_{peers}=10).

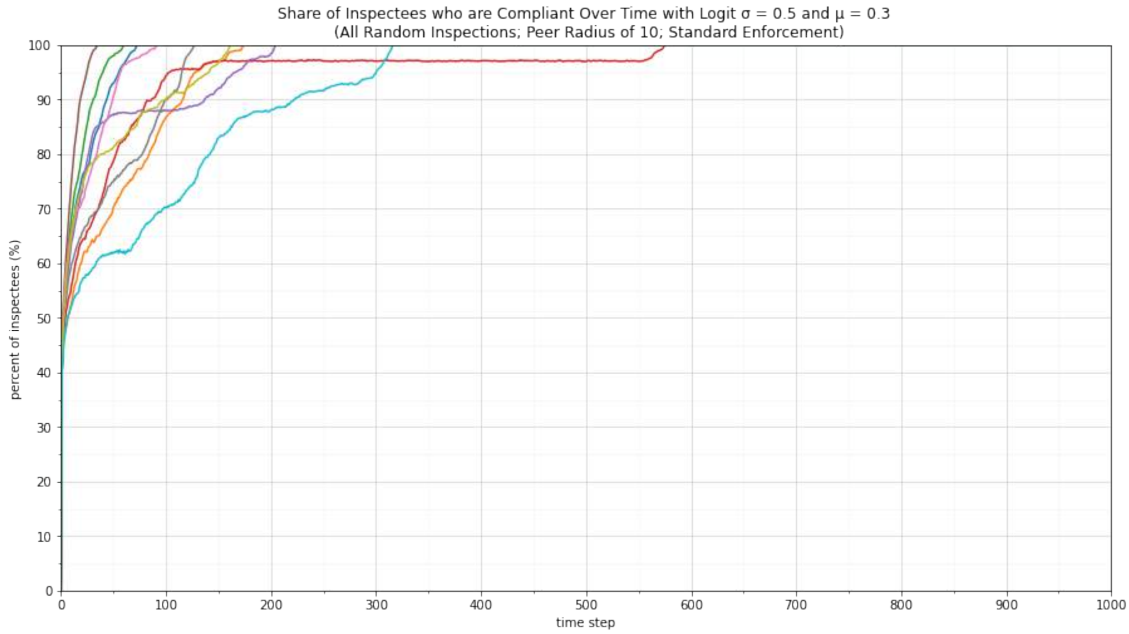

(c) σ = 0.3

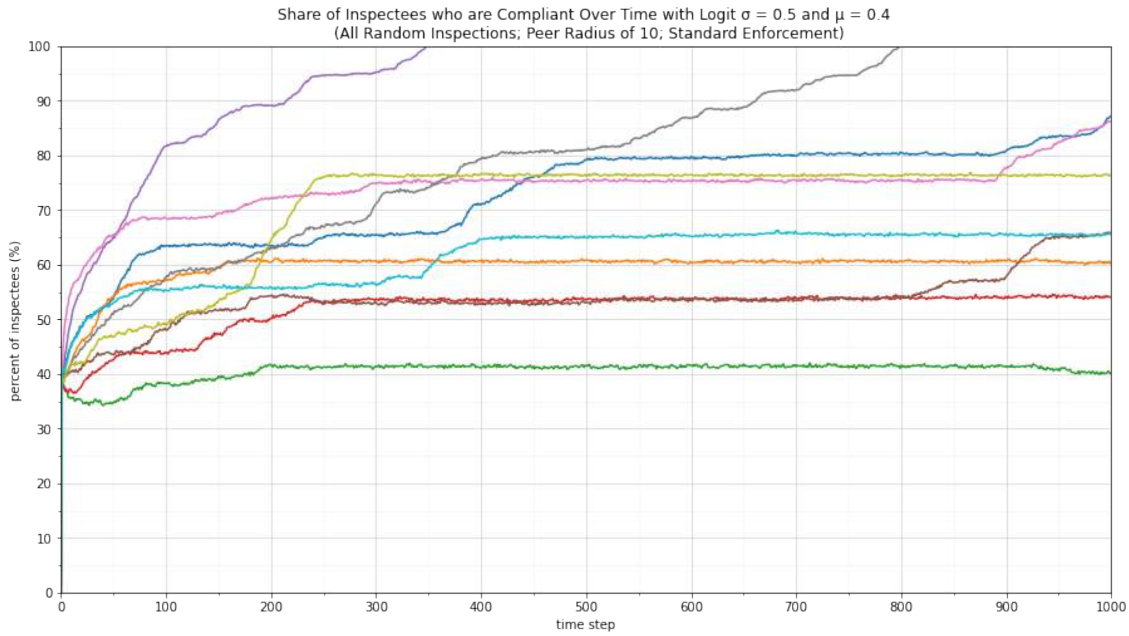

(d) σ = 0.4

Figure A.57: Average compliance rate with varying σ & μ=1 (SE, AR, r_{peers}=10, continued).

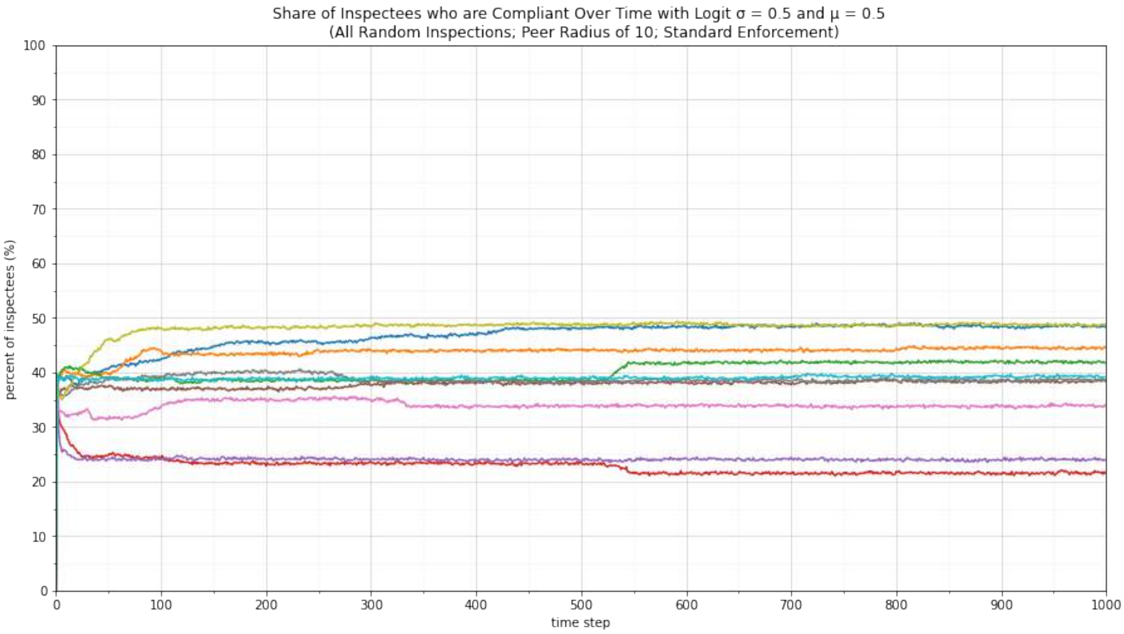

(e) $\sigma = 0.5$

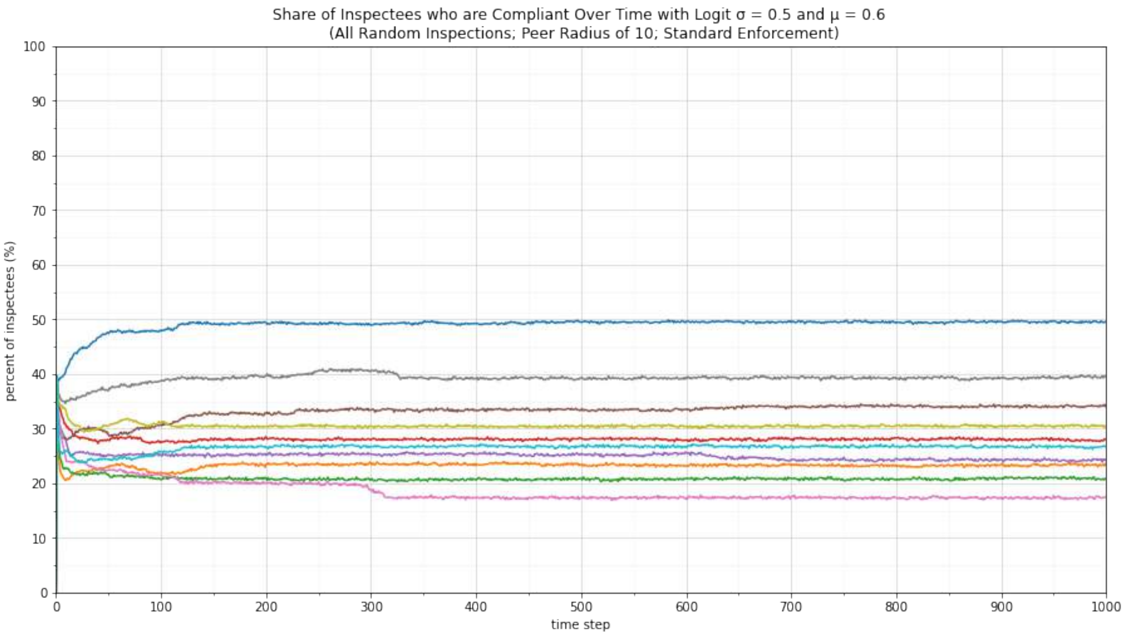

(f) $\sigma = 0.6$

Figure A.57: Average compliance rate with varying σ & μ=1 (SE, AR, r_{peers}=10, continued).

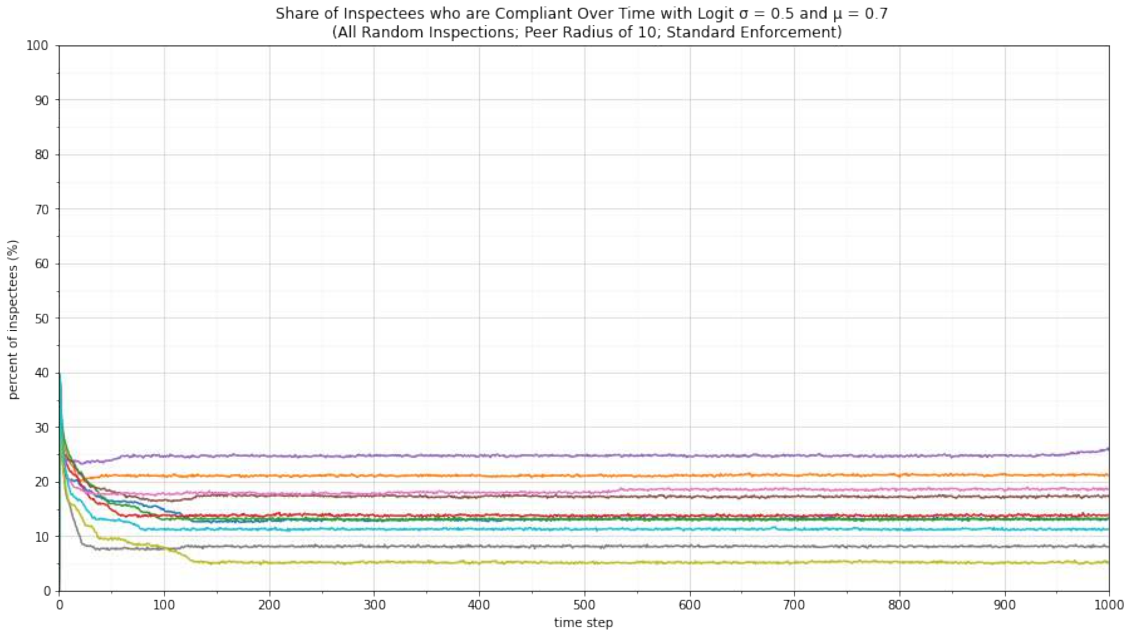

(g) $\sigma = 0.7$

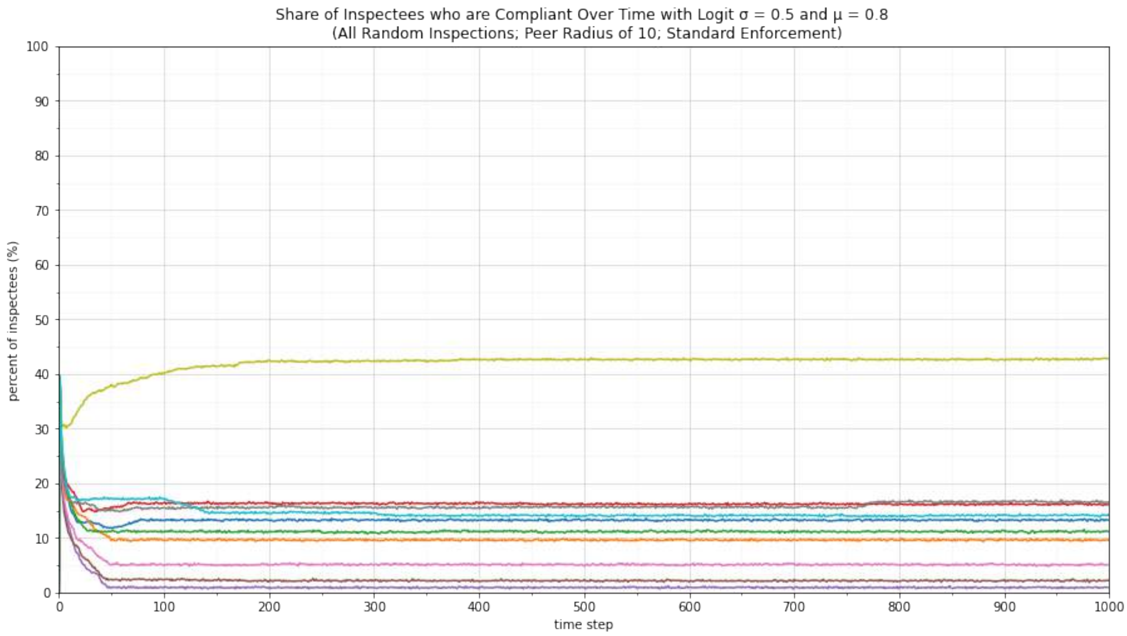

(h) $\sigma = 0.8$

Figure A.57: Average compliance rate with varying σ & μ=1 (SE, AR, r_{peers}=10, continued).

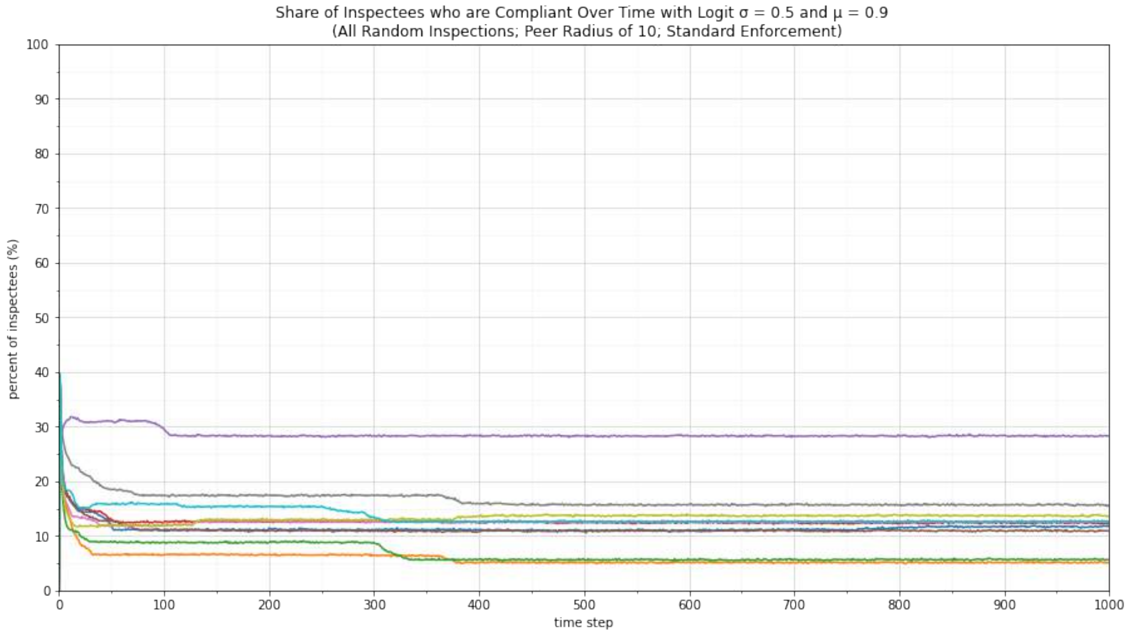

(i) σ = 0.9

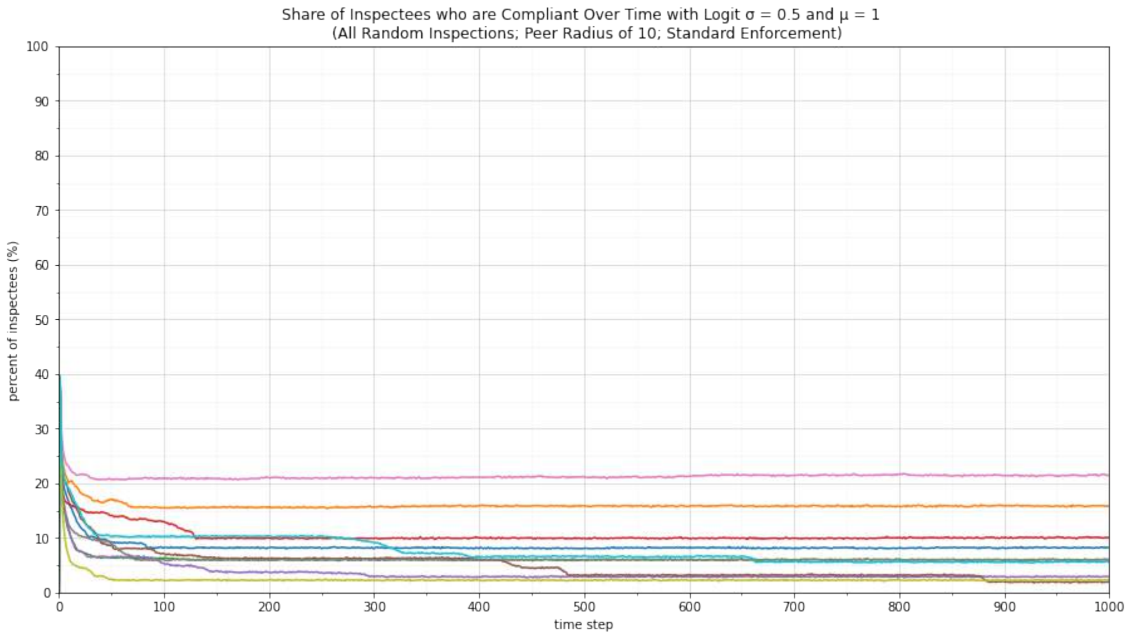

(j) σ = 1

Figure A.57: Average compliance rate with varying σ & μ=1 (SE, AR, r_{peers}=10, continued).

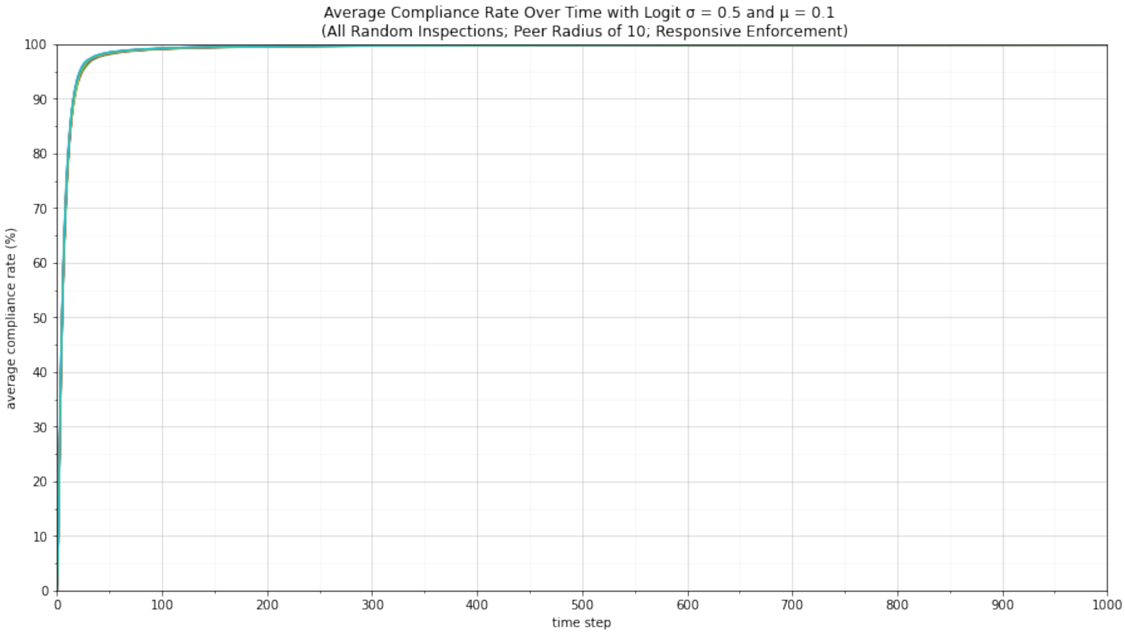

(a) σ = 0.1

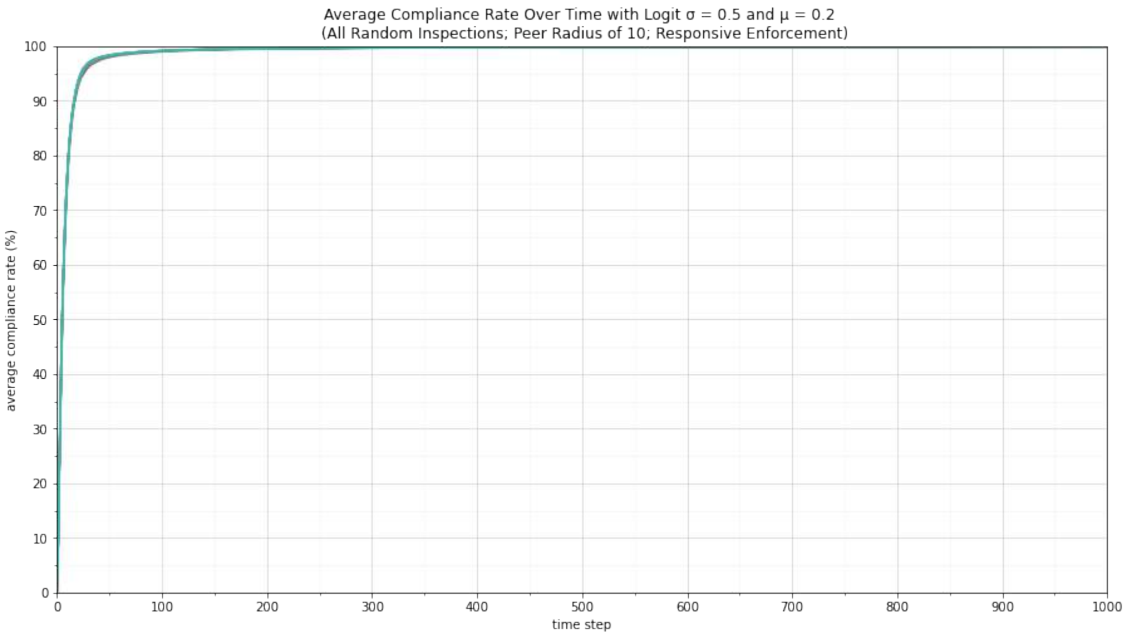

(b) σ = 0.2

Figure A.58: Average compliance rate with varying σ & μ=1 (RE, AR, r_{peers}=10).

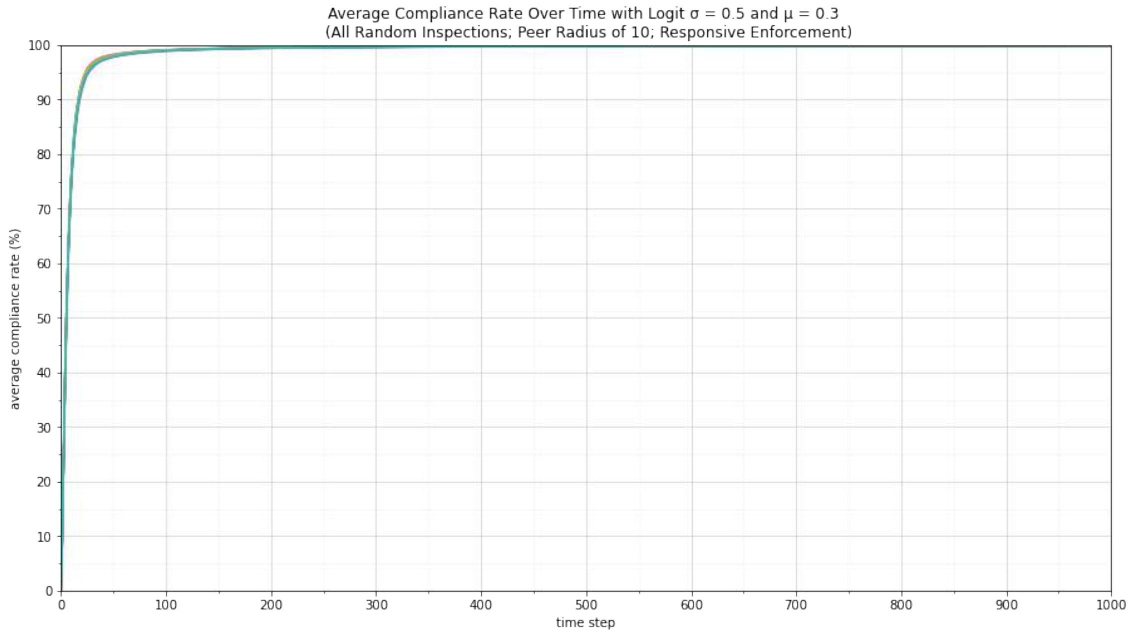

(c) σ = 0.3

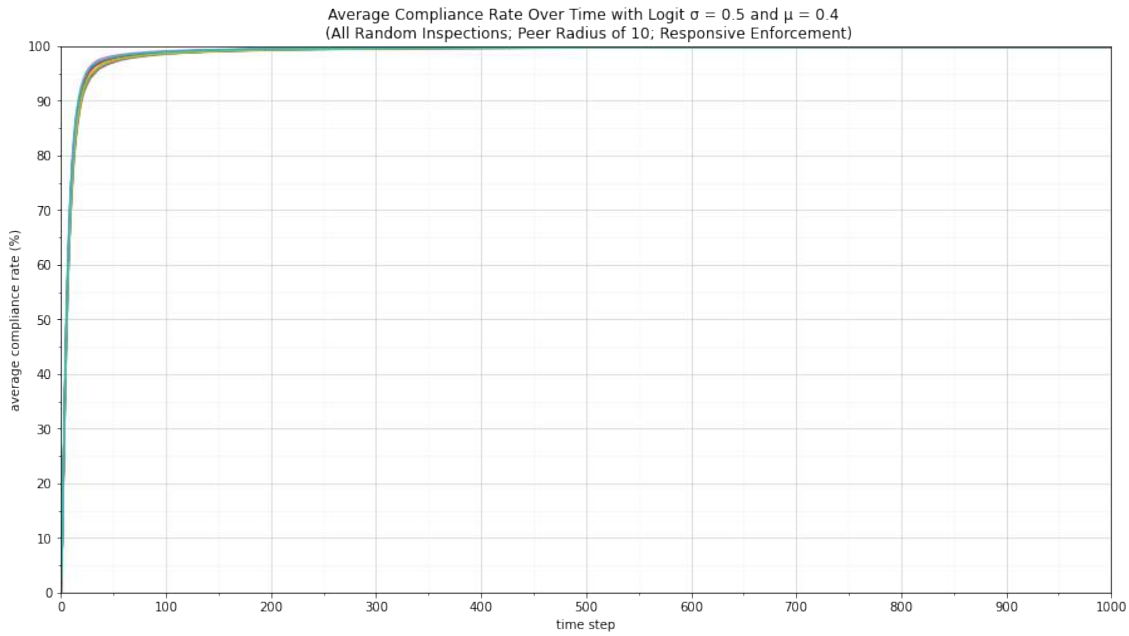

(d) σ = 0.4

Figure A.58: Average compliance rate with varying σ & μ=1 (RE, AR, r_{peers}=10, continued).

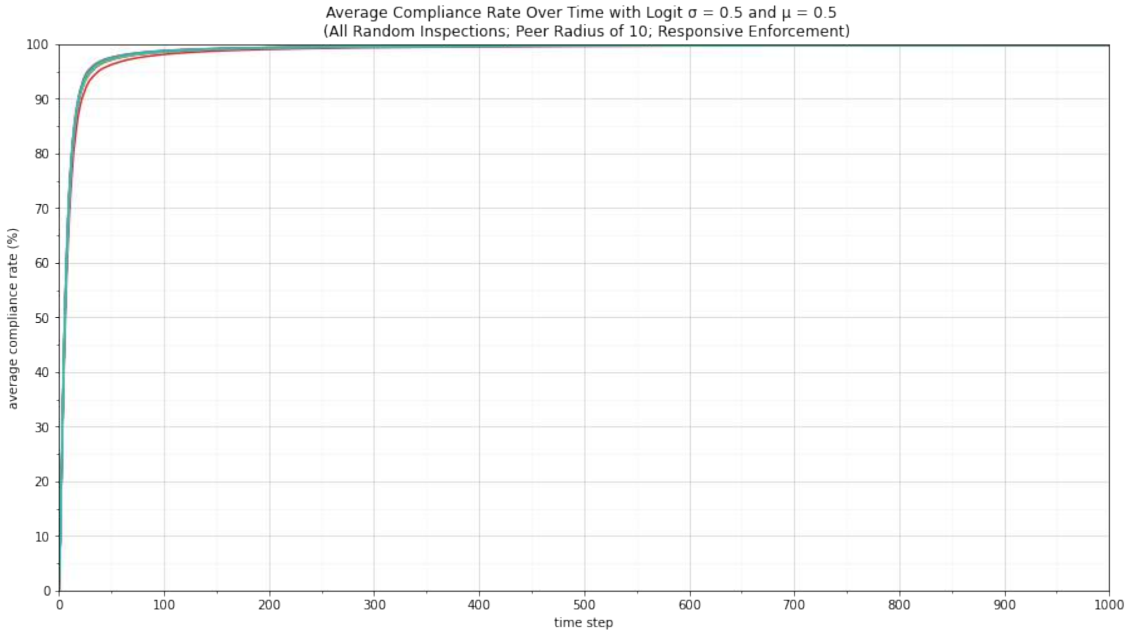

(e) σ = 0.5

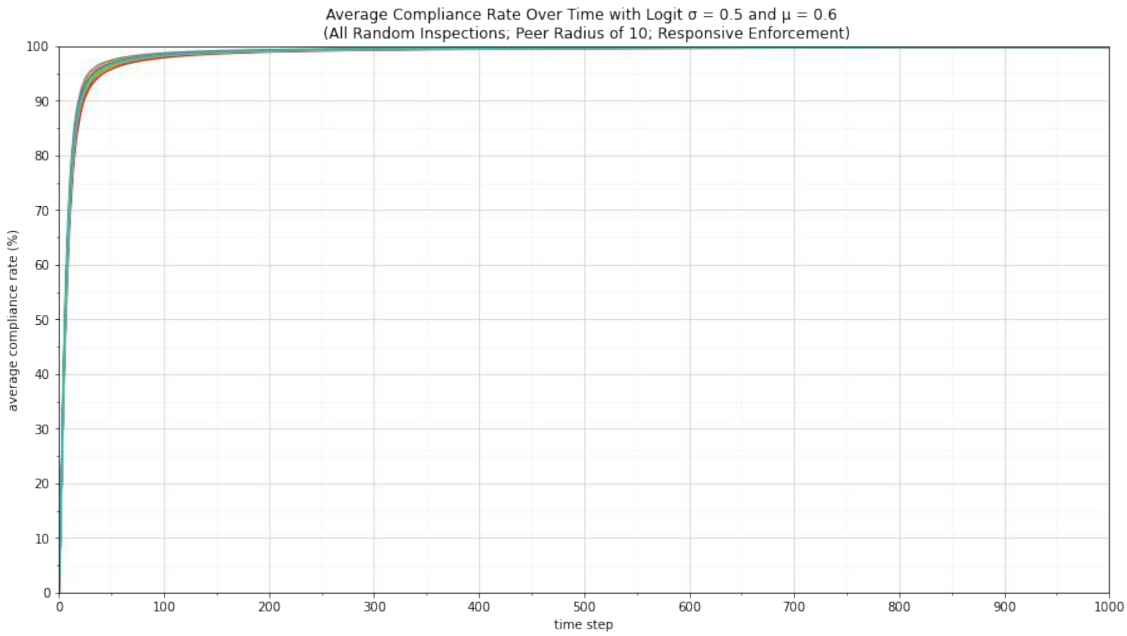

(f) σ = 0.6

Figure A.58: Average compliance rate with varying σ & μ=1 (RE, AR, r_{peers}=10, continued).

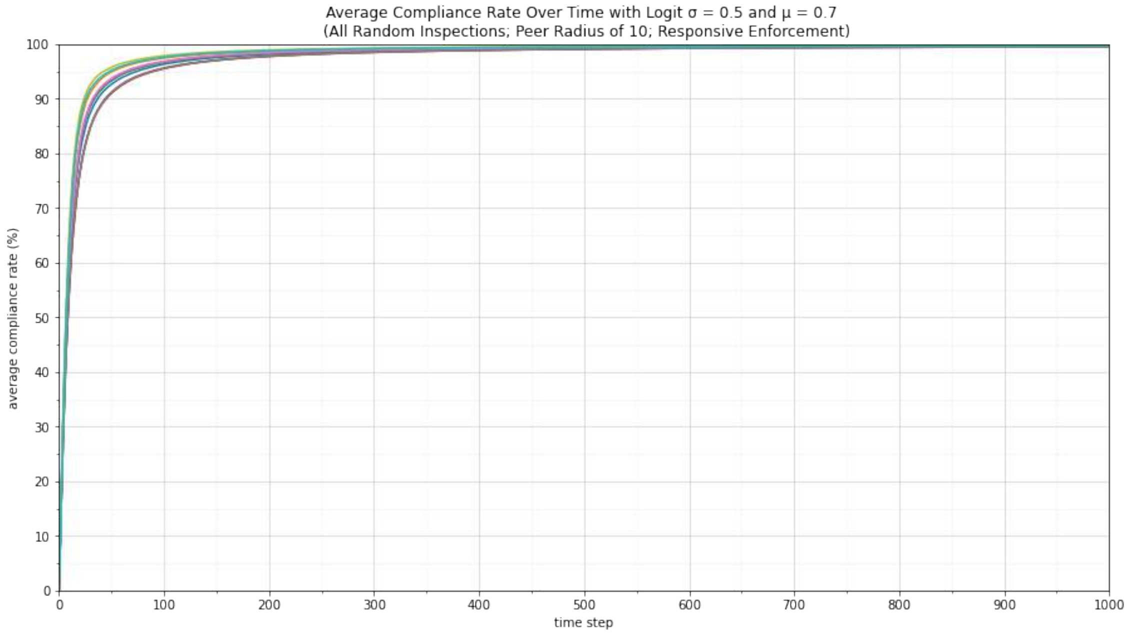

(g) $\sigma = 0.7$

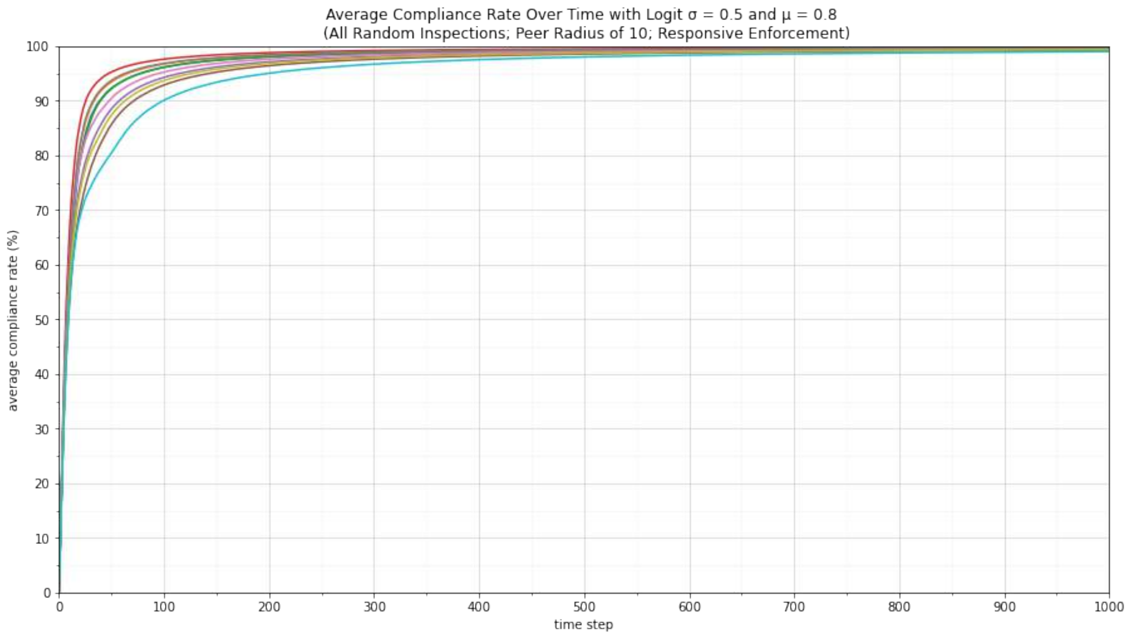

(h) $\sigma = 0.8$

Figure A.58: Average compliance rate with varying σ & μ=1 (RE, AR, r_{peers}=10, continued).

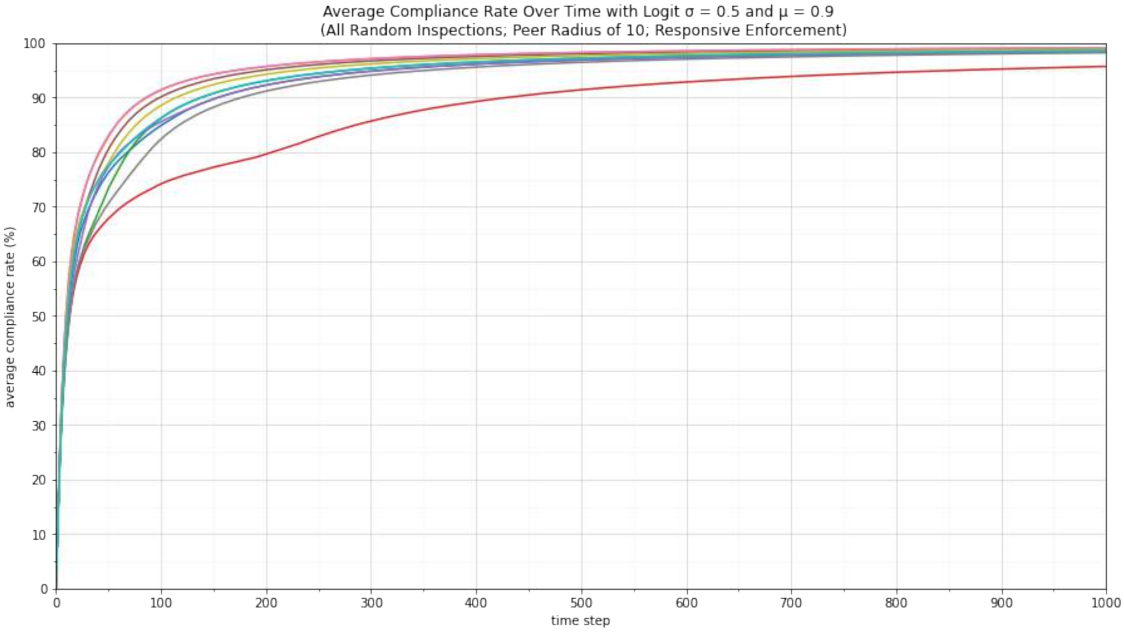

(i) σ = 0.9

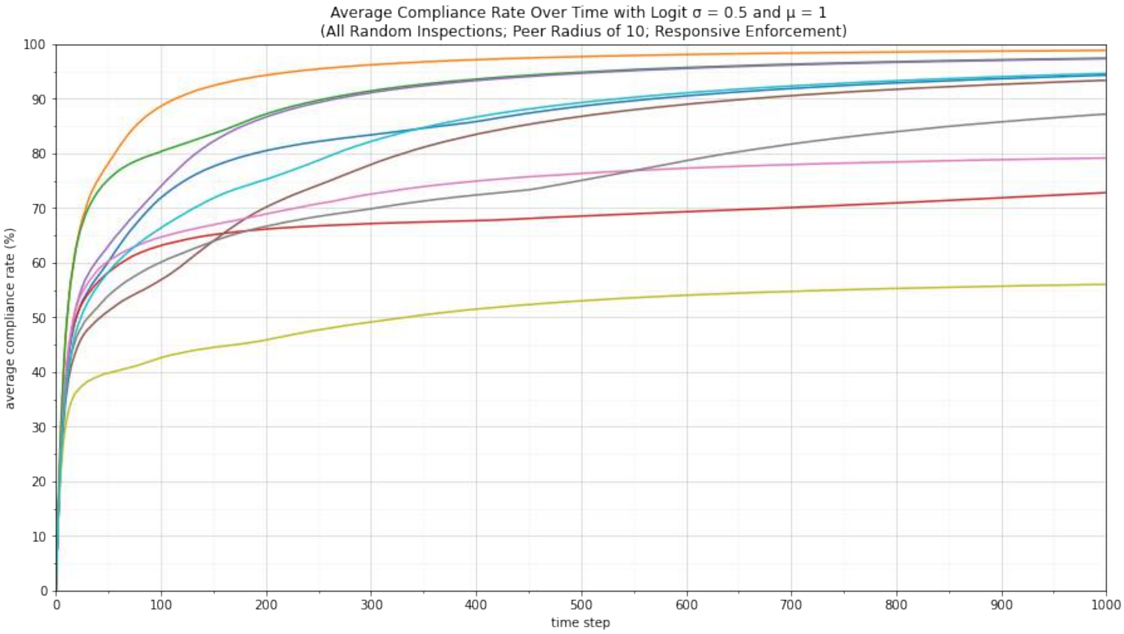

(j) σ = 1

Figure A.58: Average compliance rate with varying σ & μ=1 (RE, AR, r_{peers}=10, continued).

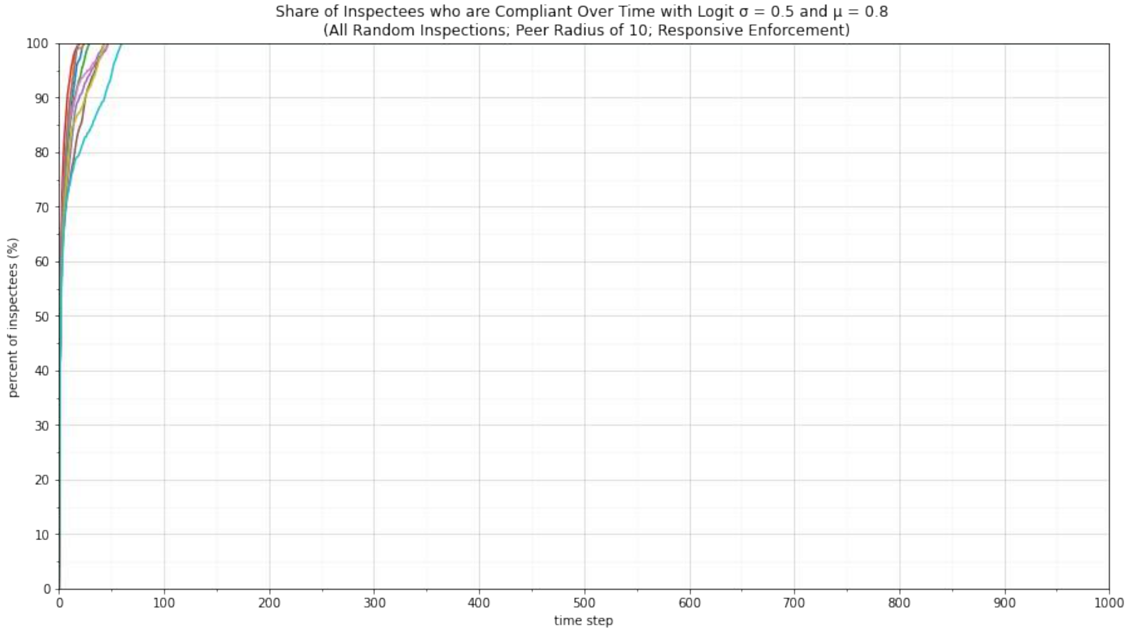

(a) σ = 0.8

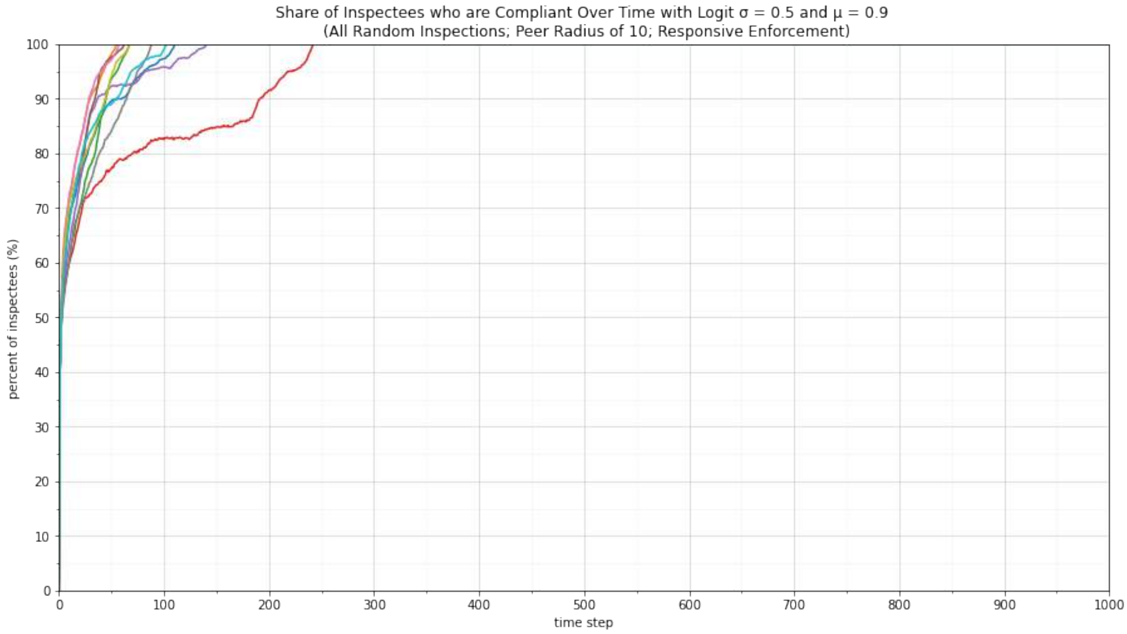

(b) σ = 0.9

Figure A.59: Share of compliant inspectees with varying σ & μ=1 (RE, AR, r_{peers}=10).
Note: Only structurally valid values of σ are shown.

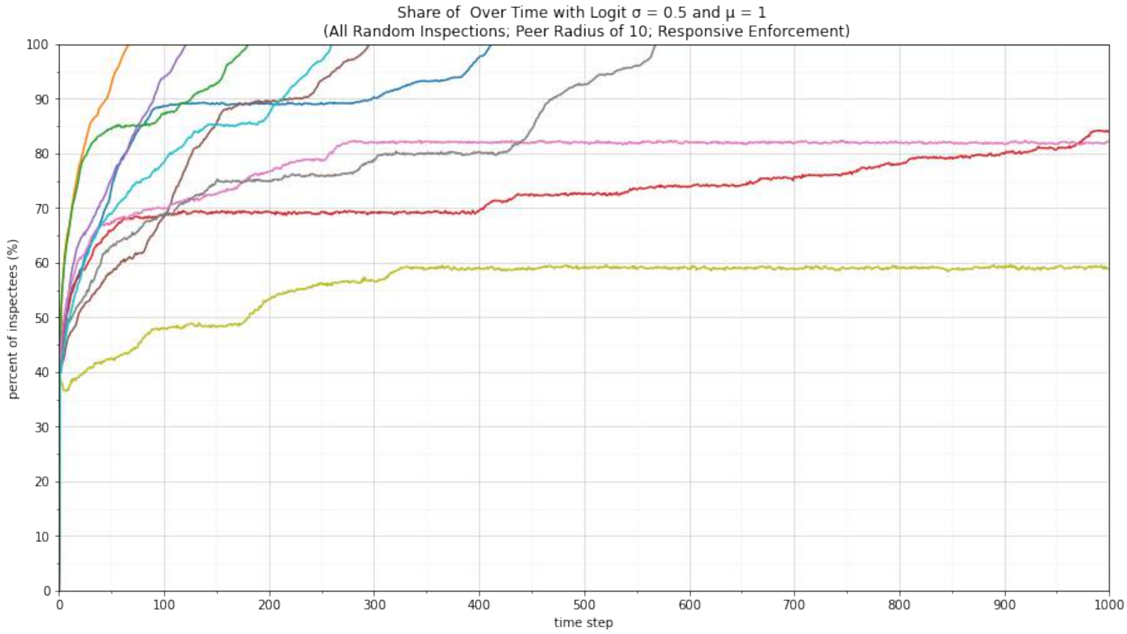

(c) $\sigma = 1$

Figure A.59: Non-compliant inspectee breakdown with varying σ & μ=1 (SE, AR, r_{peers}=10) (continued).

A.1.4 Varying Radius of Peers (r_{peers})

Standard Enforcement

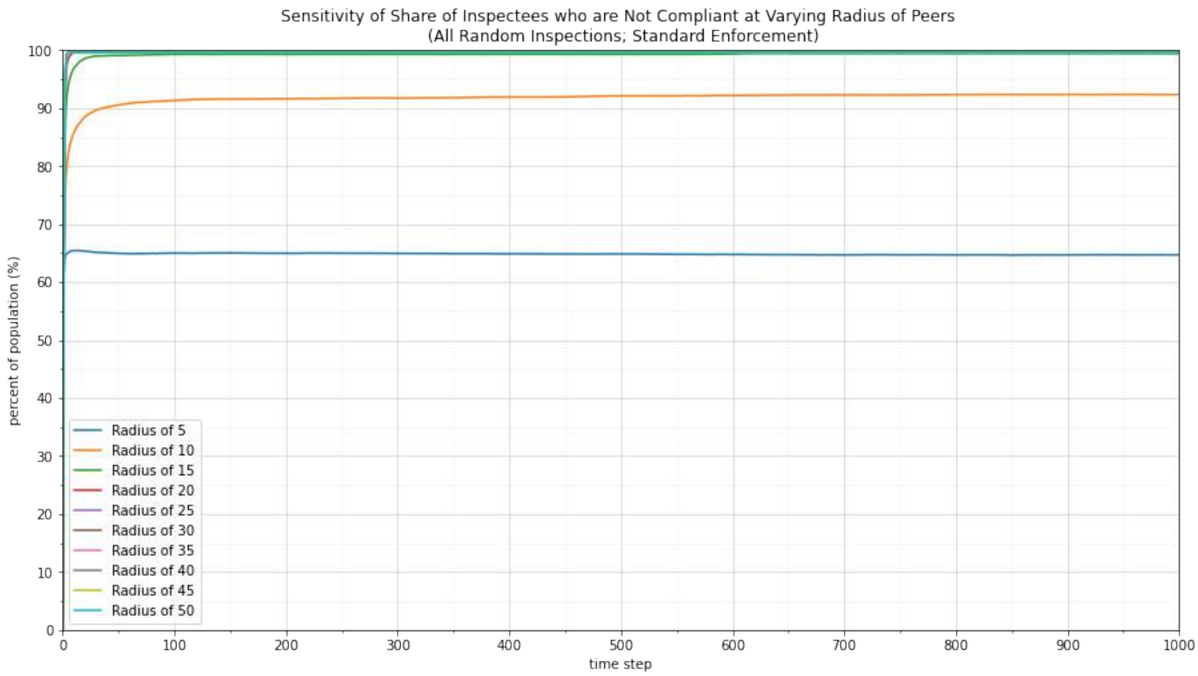

Figure A.60: Share of non-compliant inspectees with varying r_{peers} (SE).
For a detailed breakdown by compliance level, see Figure A.61.

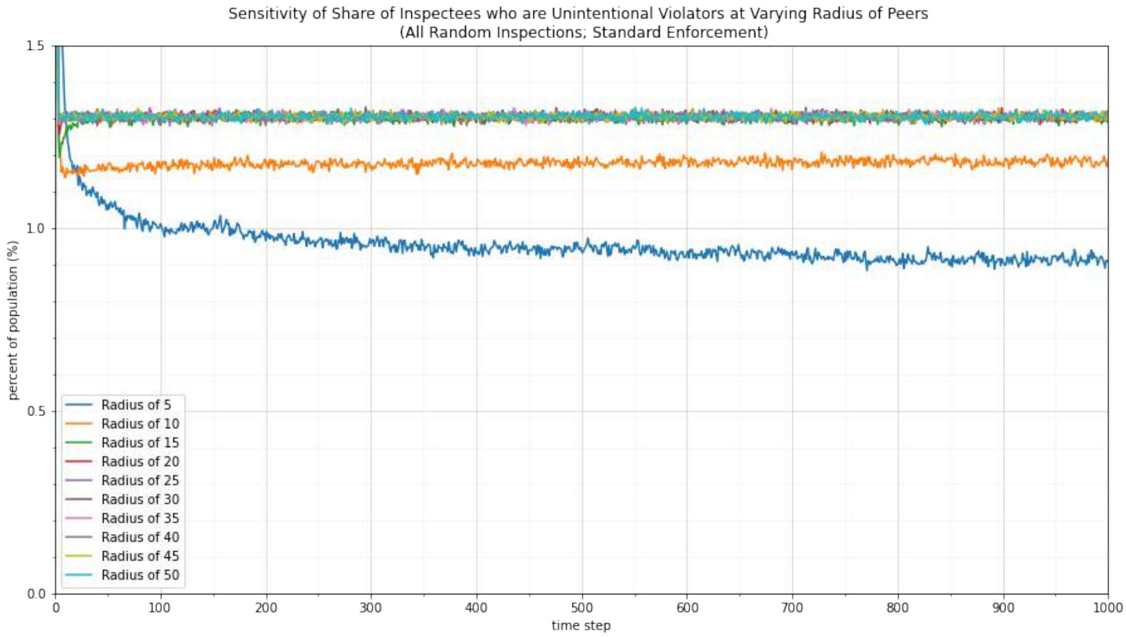

(a) Unintentional violators.

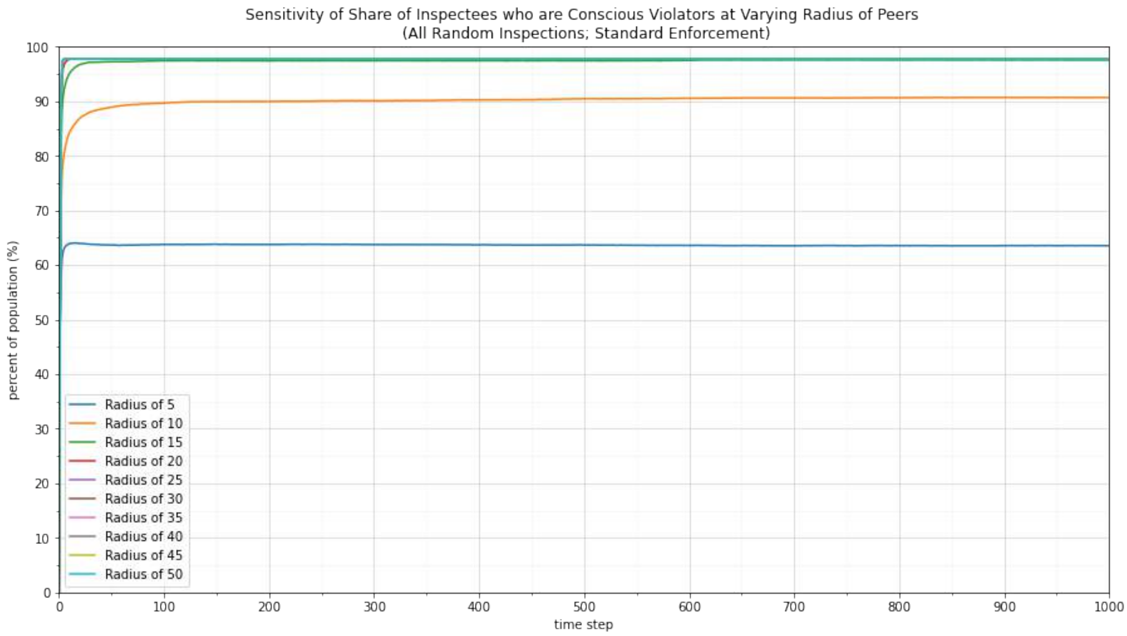

(b) Conscious violators.

Figure A.61: Non-compliant inspectee breakdown with varying r_{peers} (SE, AR).

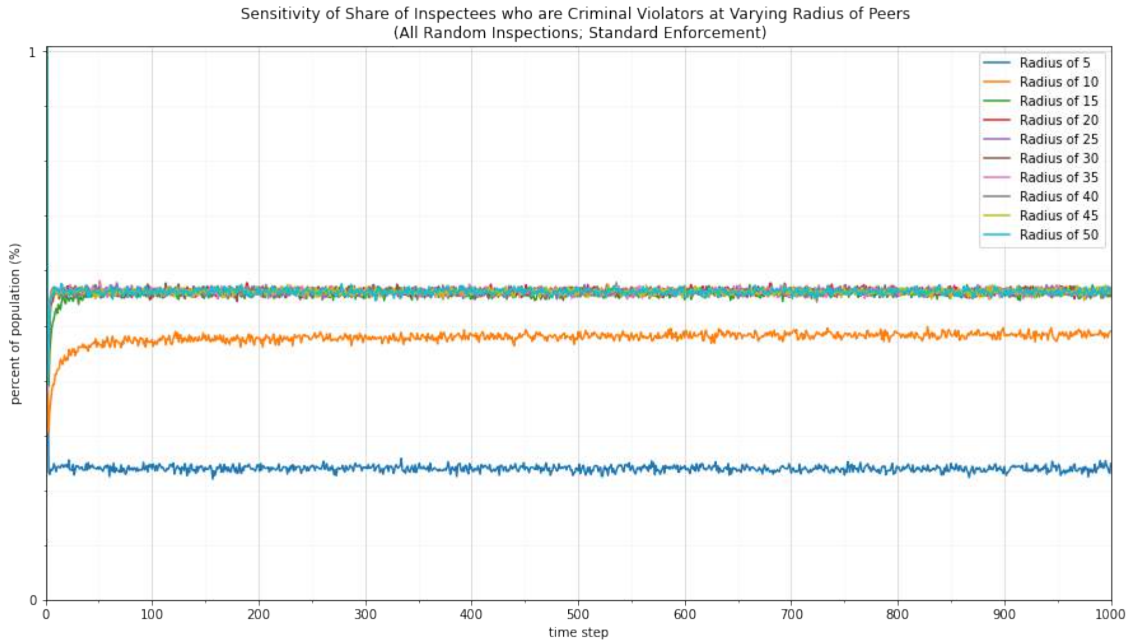

(c) Criminal violators.

Figure A.61: Non-compliant inspectee breakdown with varying r_{peers} (SE, AR, continued).

Responsive Enforcement

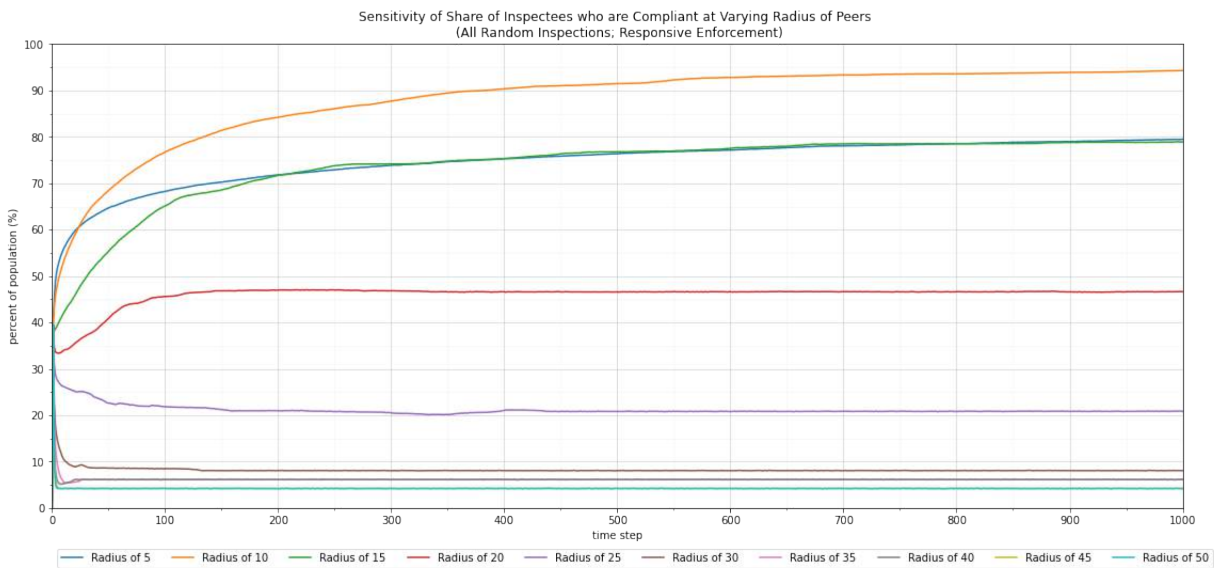

Figure A.62: Share of compliant inspectees with varying r_{peers} (RE).

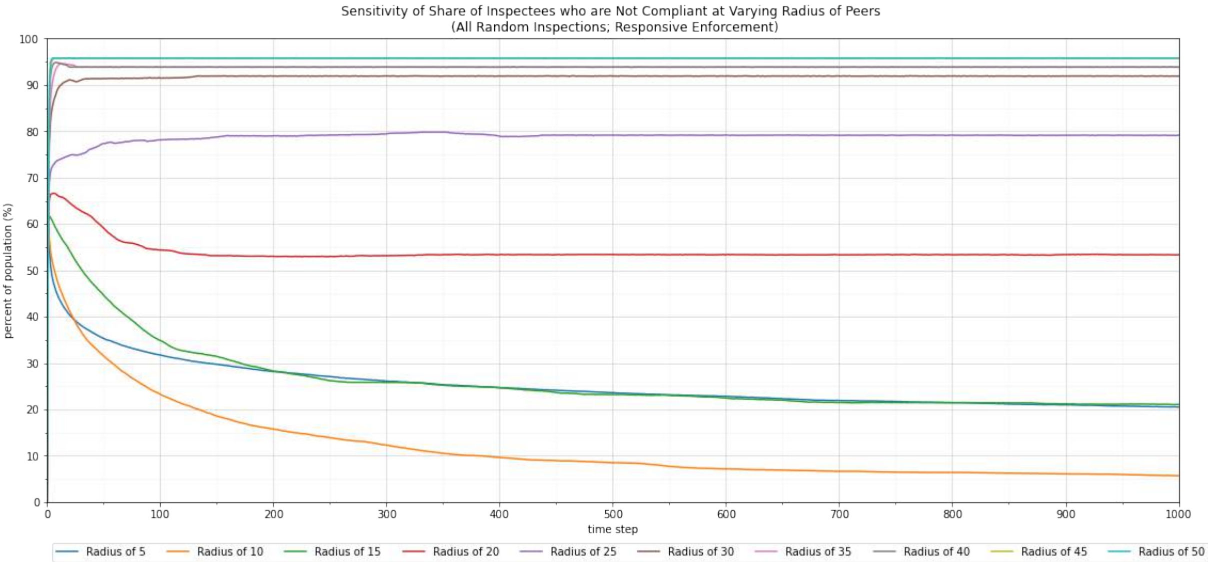

Figure A.63: Share of non-compliant inspectees with varying r_{peers} (RE).

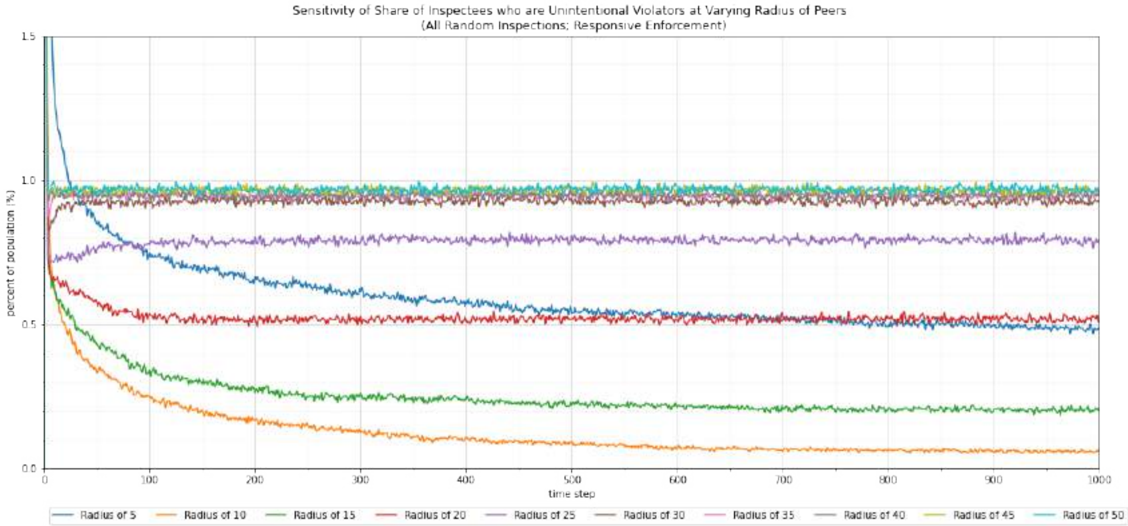

(a) Unintentional violators.

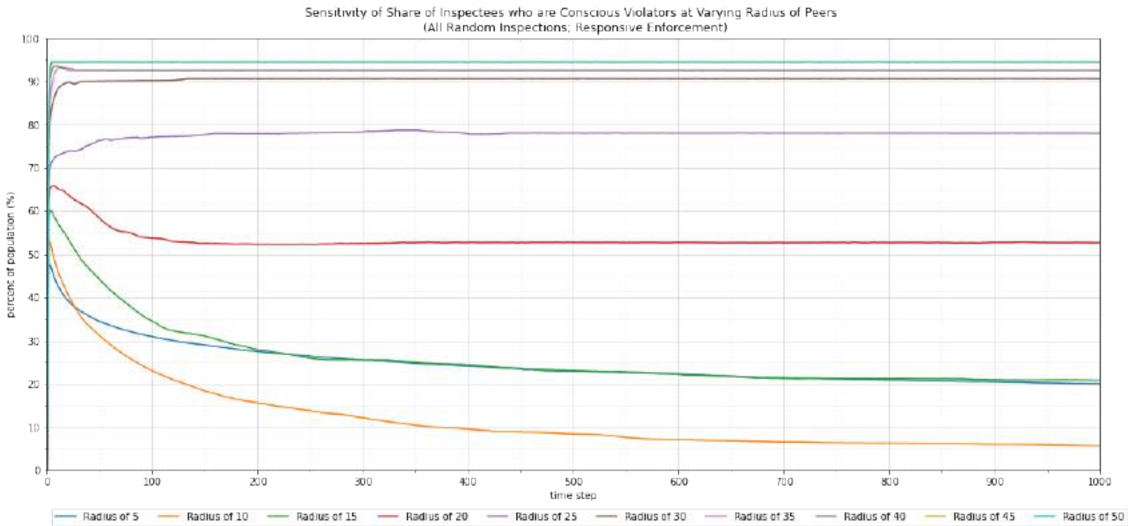

(b) Conscious violators.

Figure A.64: Non-compliant inspectee breakdown with varying r_{peers} (RE, AR).

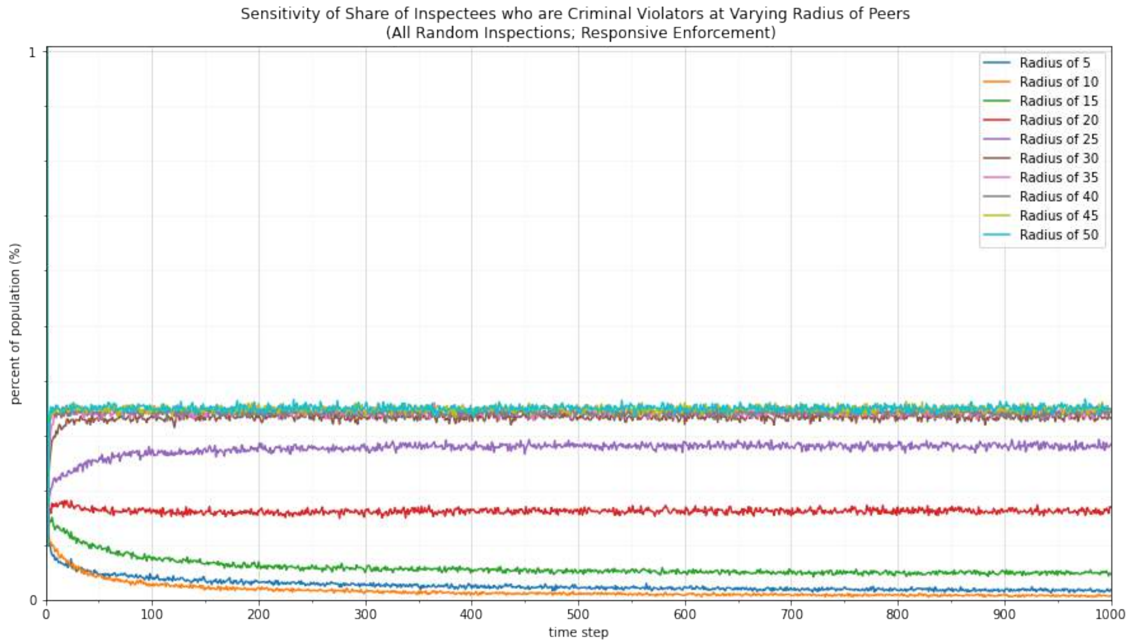

(c) Criminal violators.

Figure A.64: Non-compliant inspectee breakdown with varying r_{peers} (RE, AR, continued).

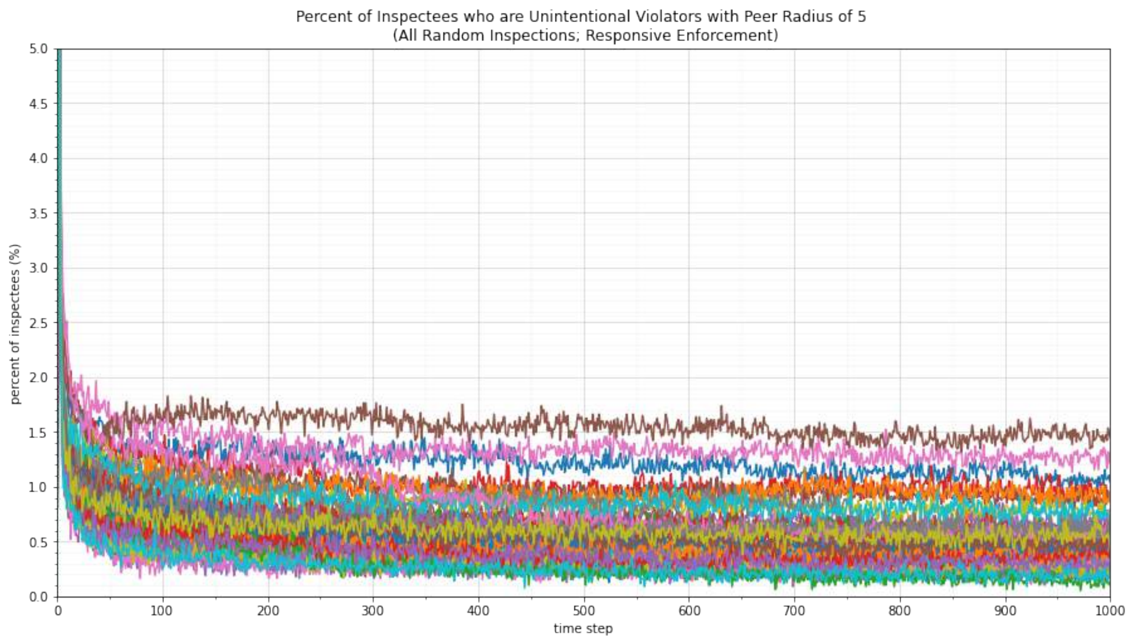

(a) Unintentional violators.

Figure A.65: Non-compliant inspectee breakdown with r_{peers}=5 (RE, AR).

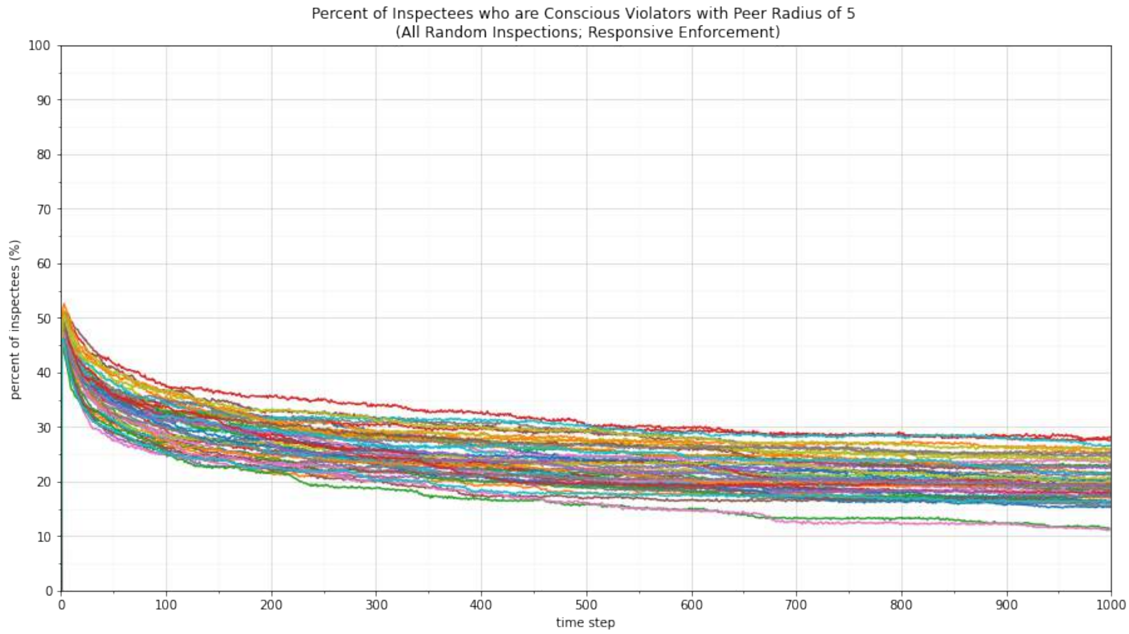

(b) Conscious violators.

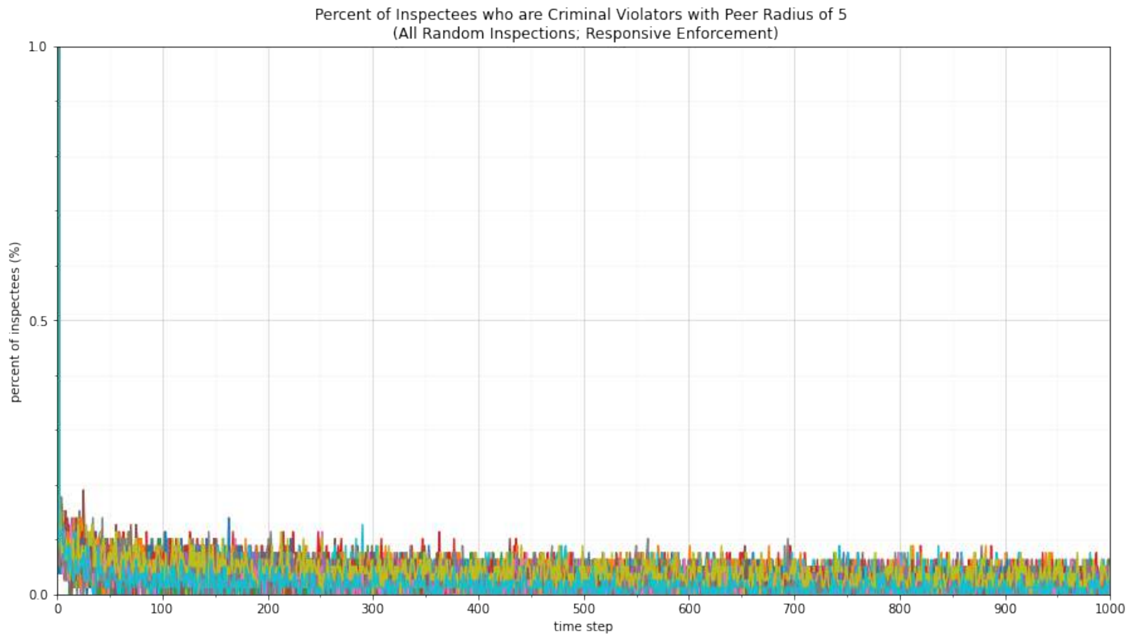

(c) Criminal violators.

Figure A.65: Non-compliant inspectee breakdown with r_{peers}=5 (RE, AR, continued).

Path Dependency

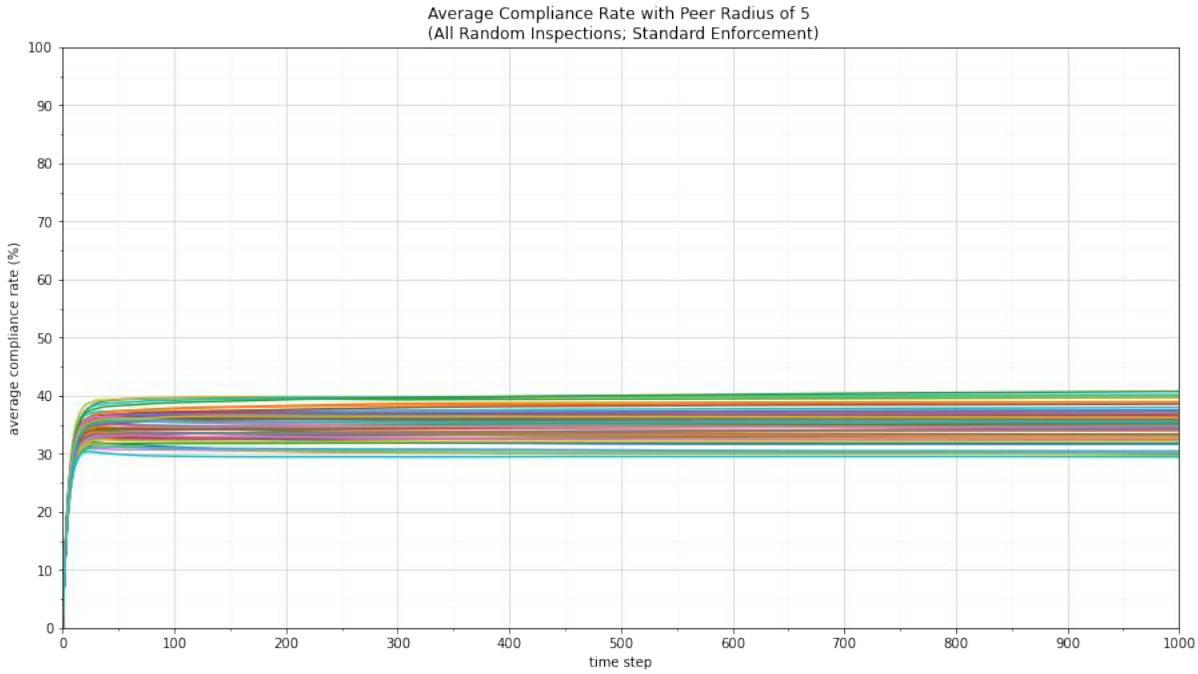

Figure A.66: Average compliance rate with r_{peers}=5 (SE).

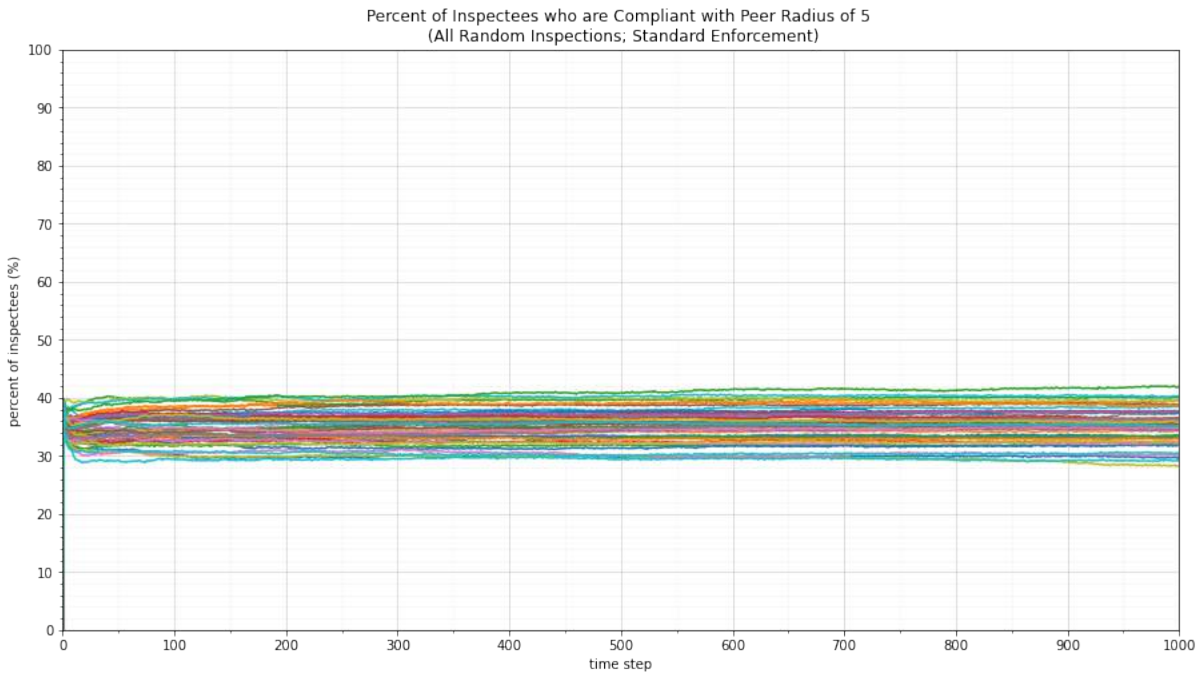

Figure A.67: Percent of compliant inspectees with r_{peers}=5 (SE).

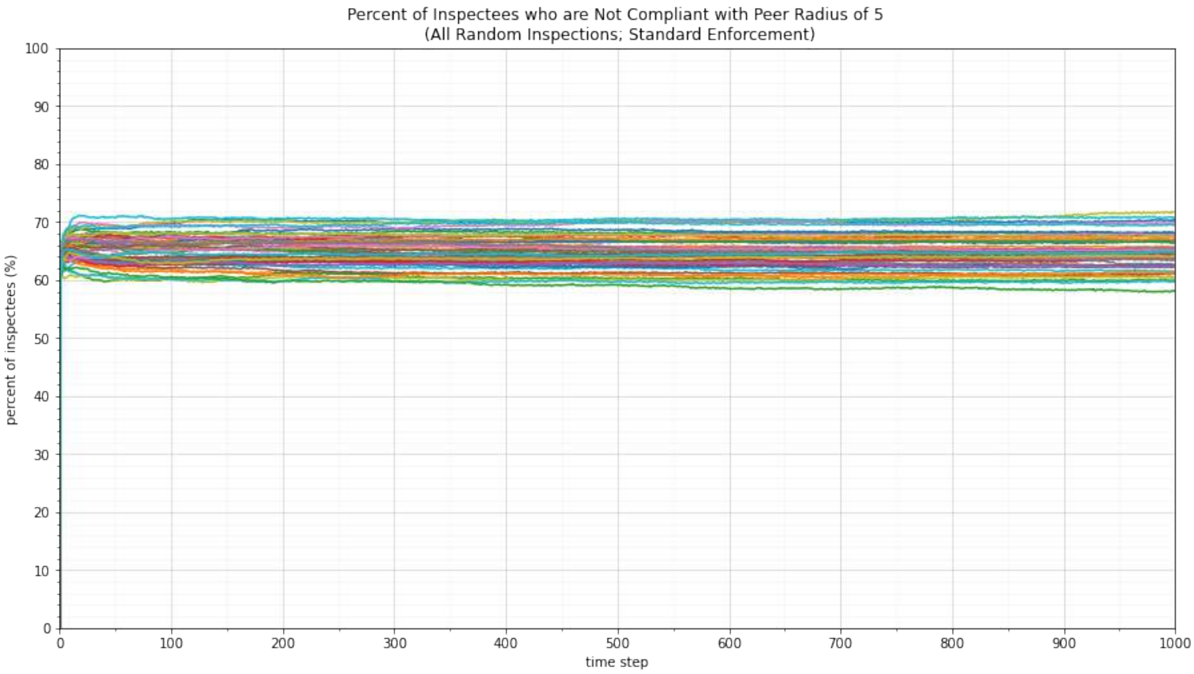

Figure A.68: Percent of non-compliant inspectees with r_{peers}=5 (SE).

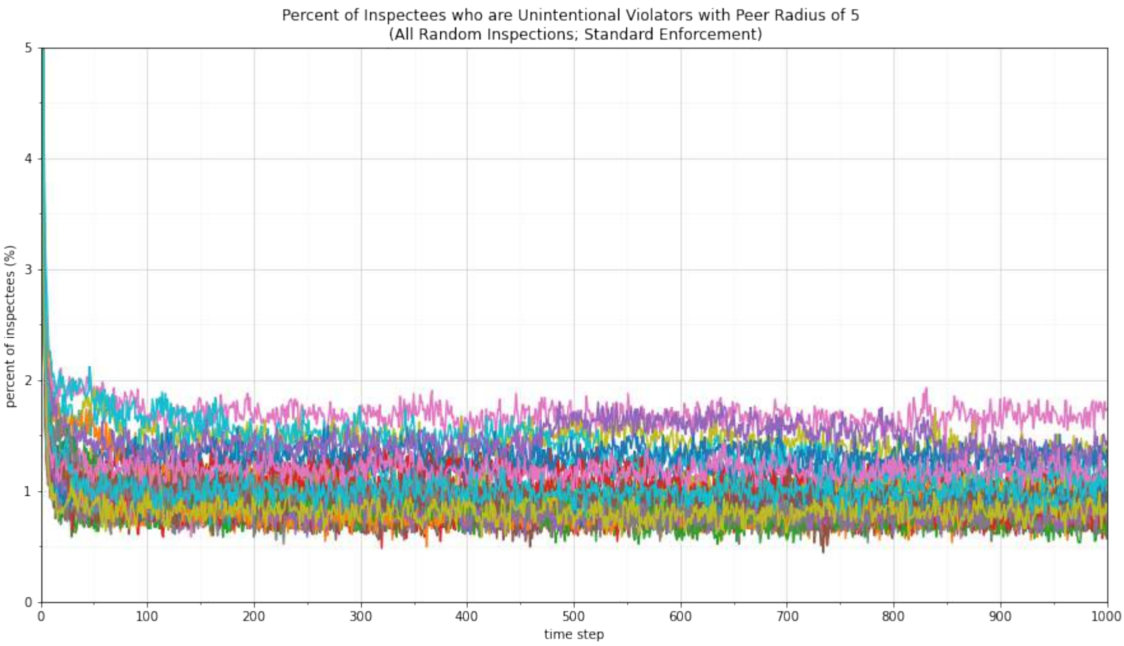

(a) Unintentional violators.

Figure A.69: Non-compliant inspectee breakdown with r_{peers}=5 (SE, AR).

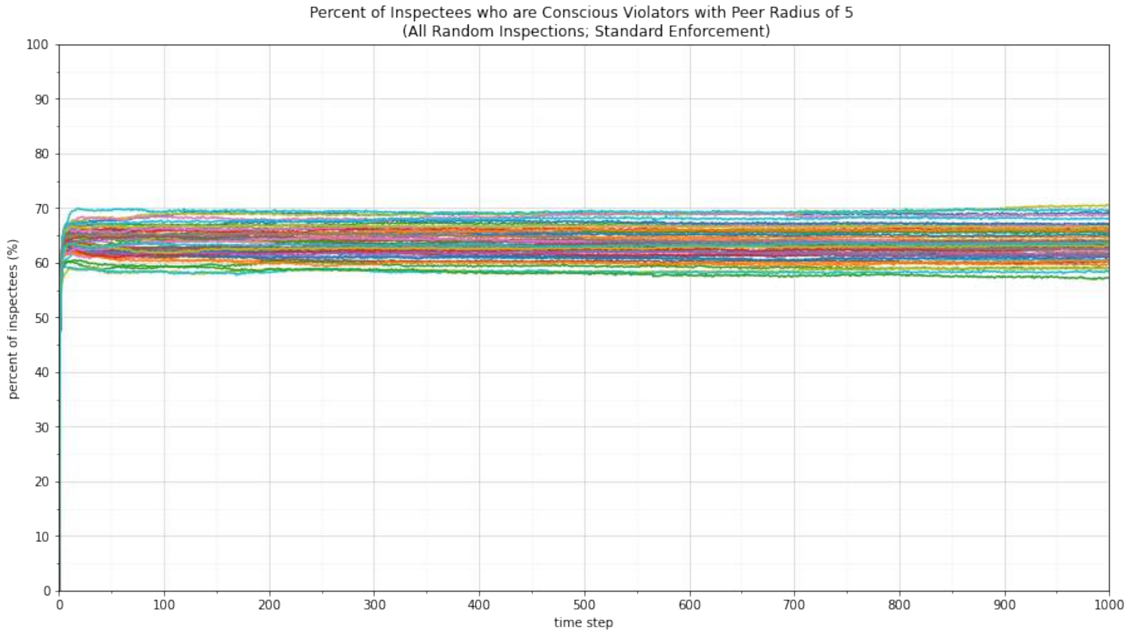

(b) Conscious violators.

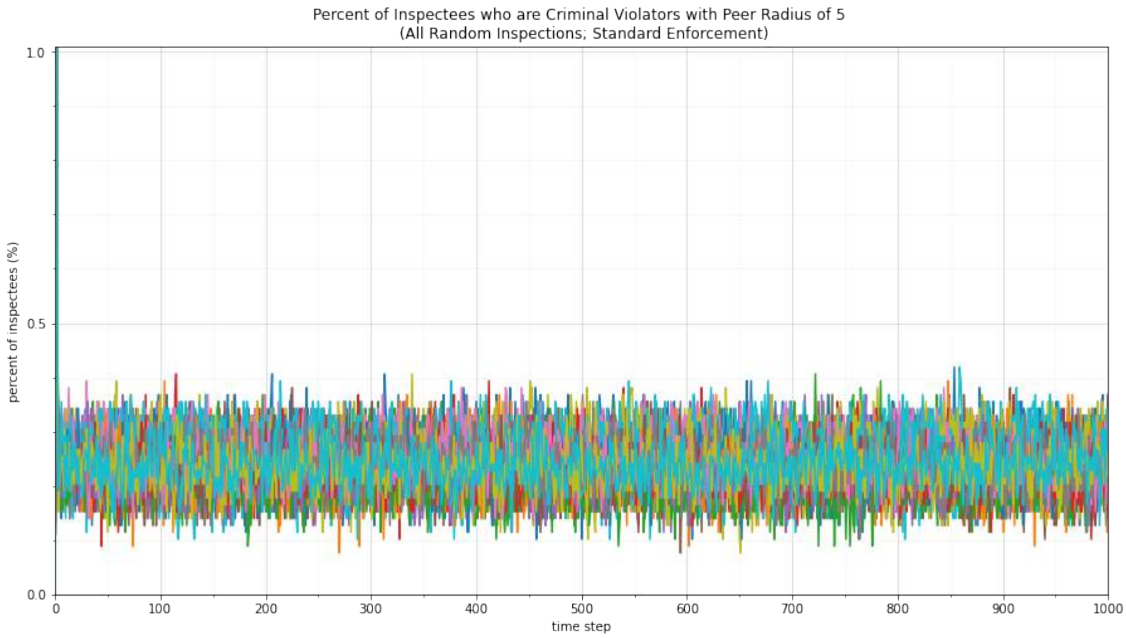

(c) Criminal violators.

Figure A.69: Non-compliant inspectee breakdown with r_{peers}=5 (SE, AR, continued).

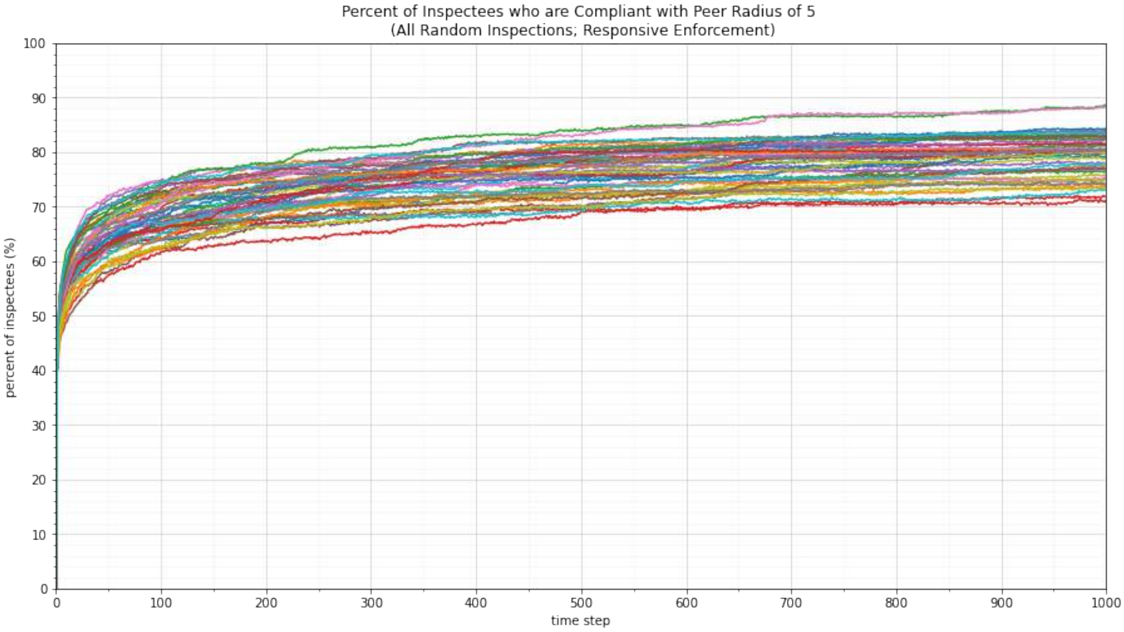

(a) Compliant inspectees.

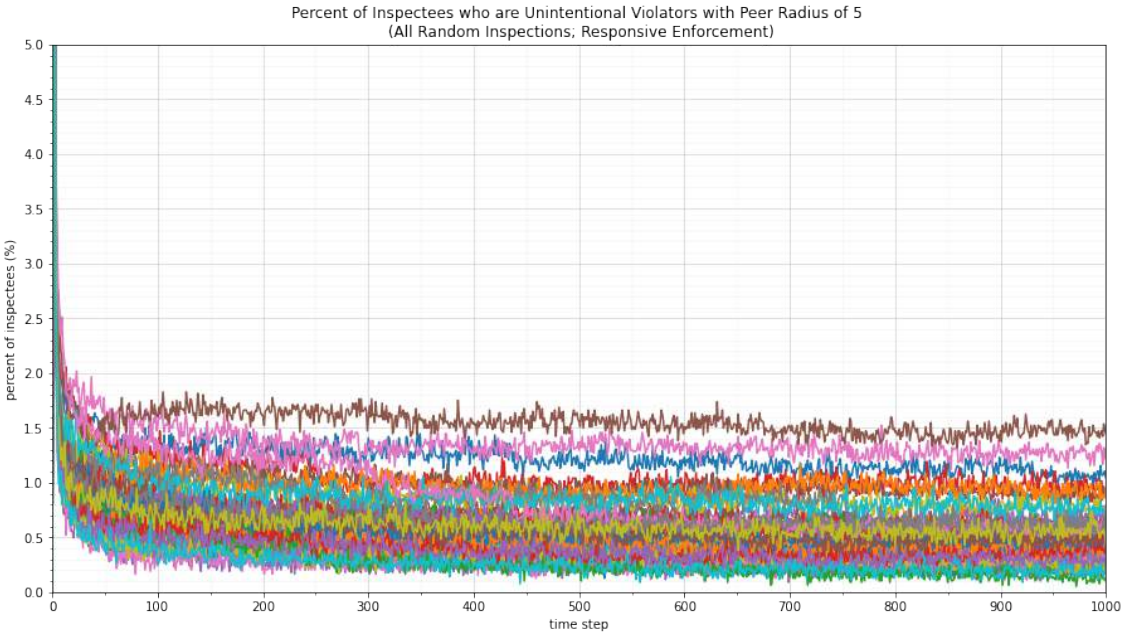

(b) Unintentional violators.

Figure A.70: Non-compliant inspectee breakdown with r_{peers}=5 (RE, AR).

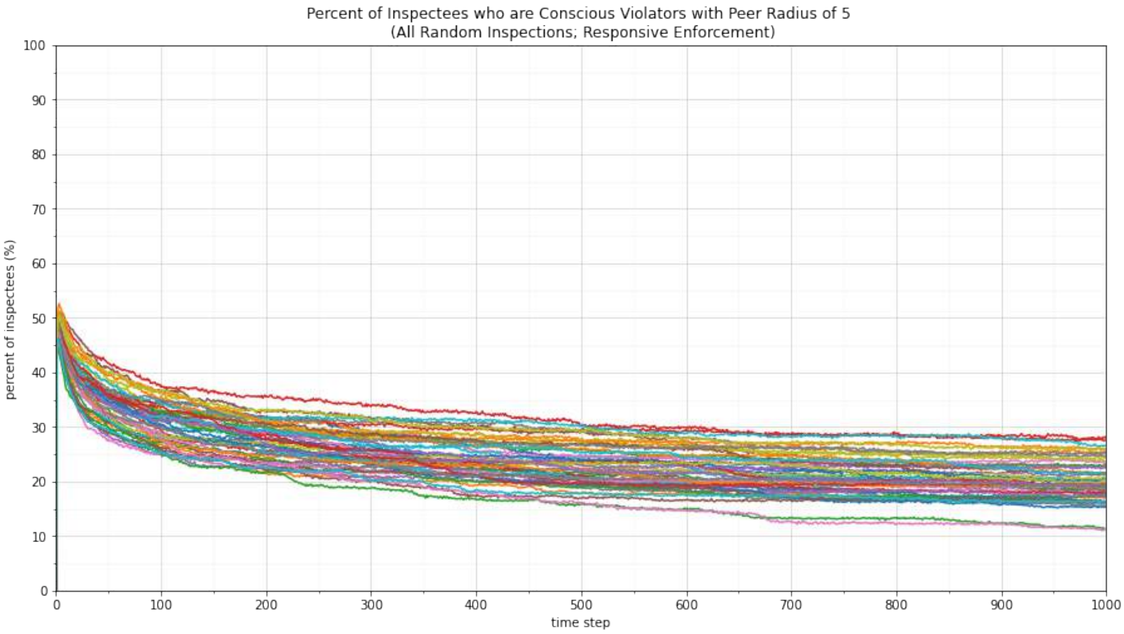

(c) Conscious violators.

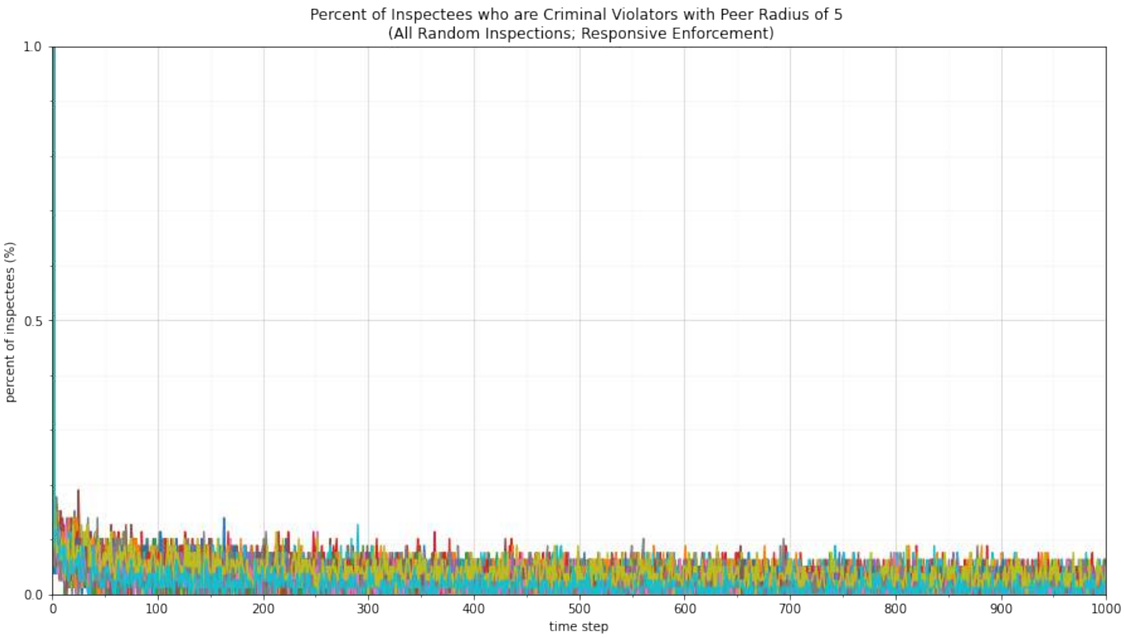

(d) Criminal violators.

Figure A.70: Non-compliant inspectee breakdown with r_{peers}=5 (RE, AR, continued).

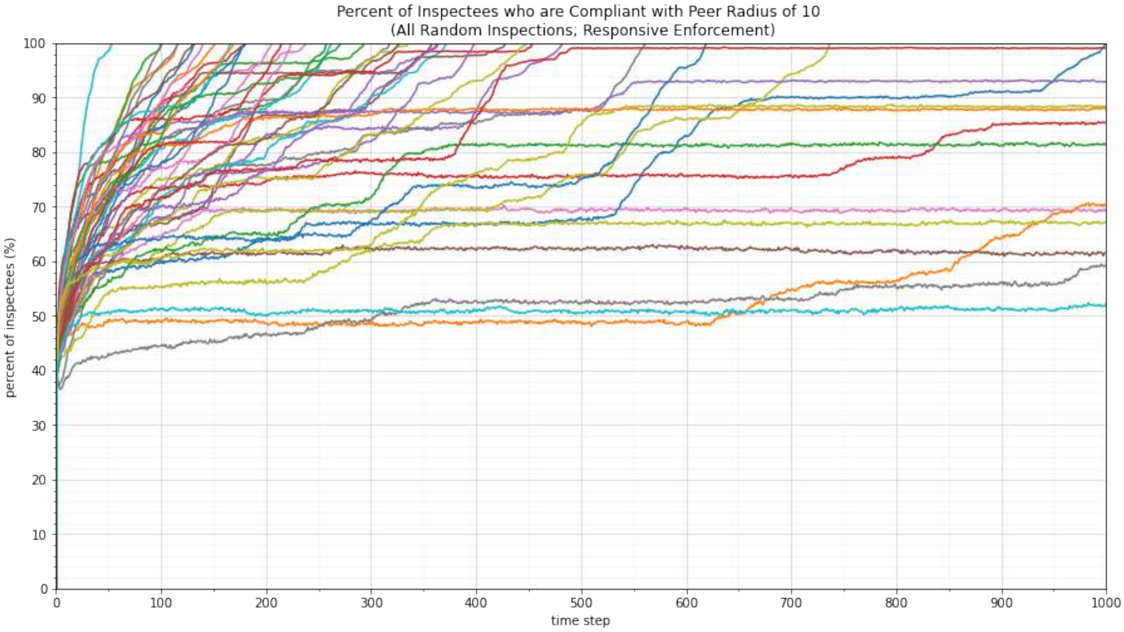

(a) Compliant inspectees.

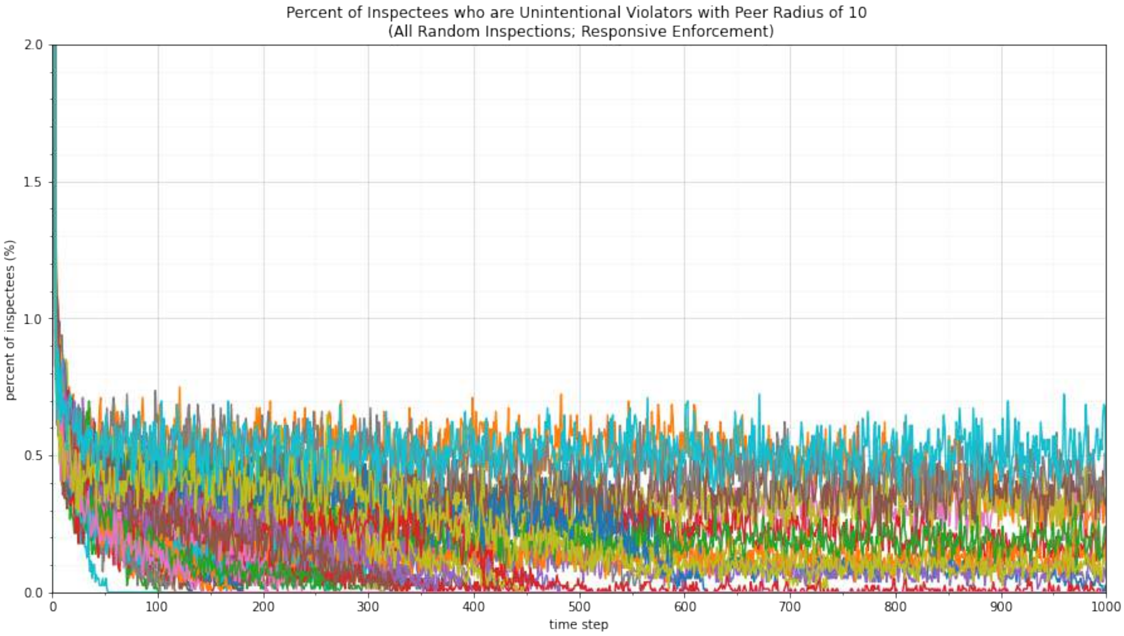

(b) Unintentional violators.

Figure A.71: Non-compliant inspectee breakdown with r_{peers}=10 (RE, AR).

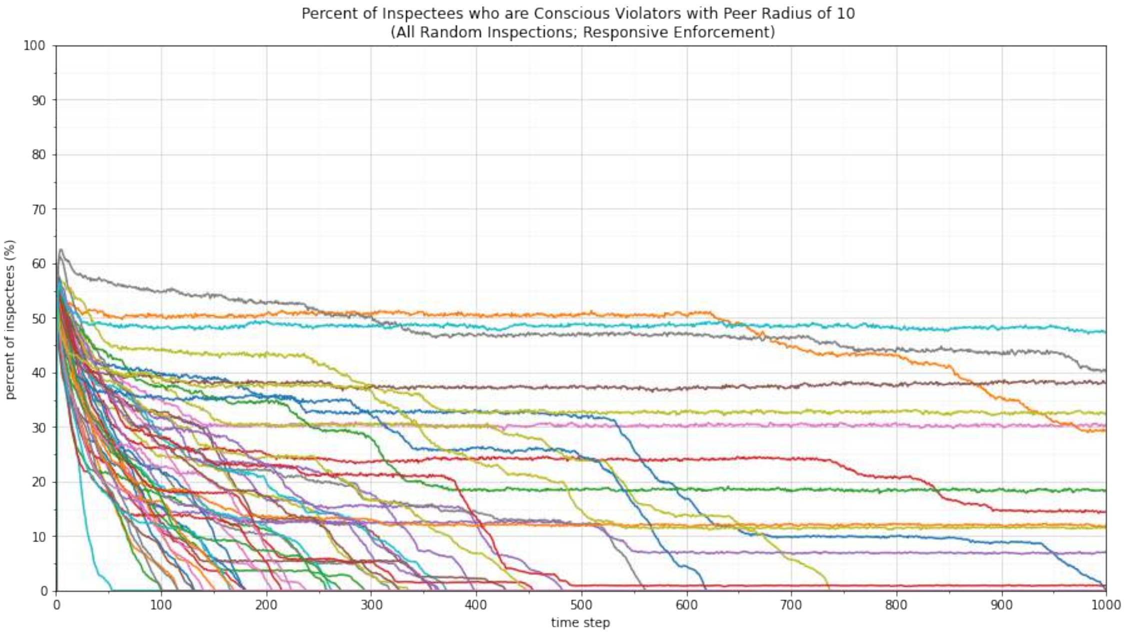

(c) Conscious violators.

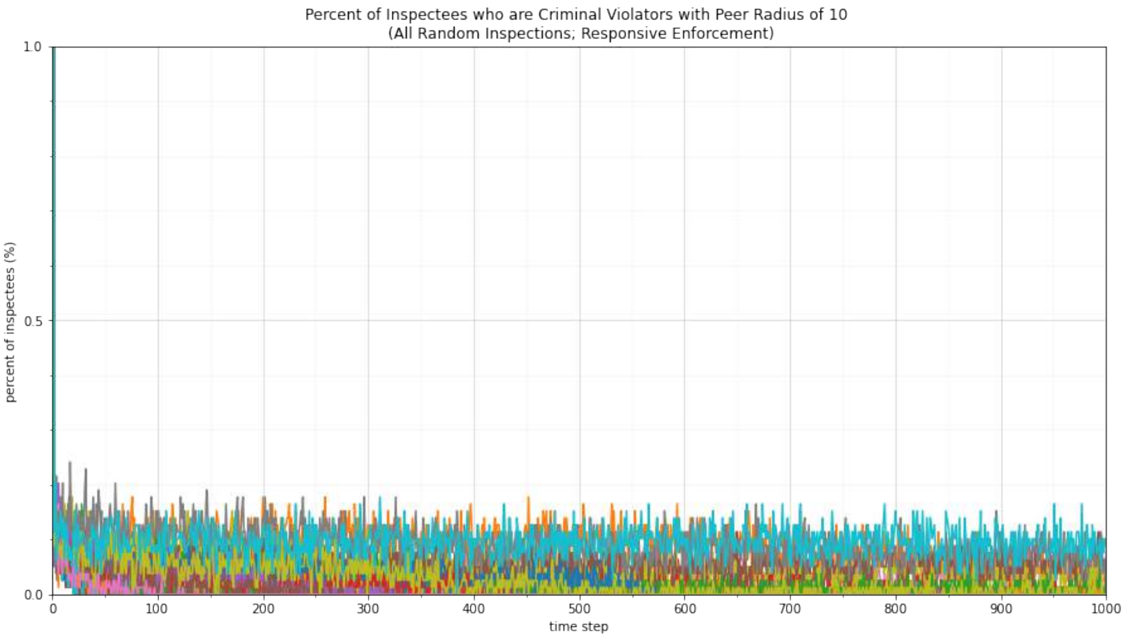

(d) Criminal violators.

Figure A.71: Non-compliant inspectee breakdown with r_{peers}=10 (RE, AR, continued).

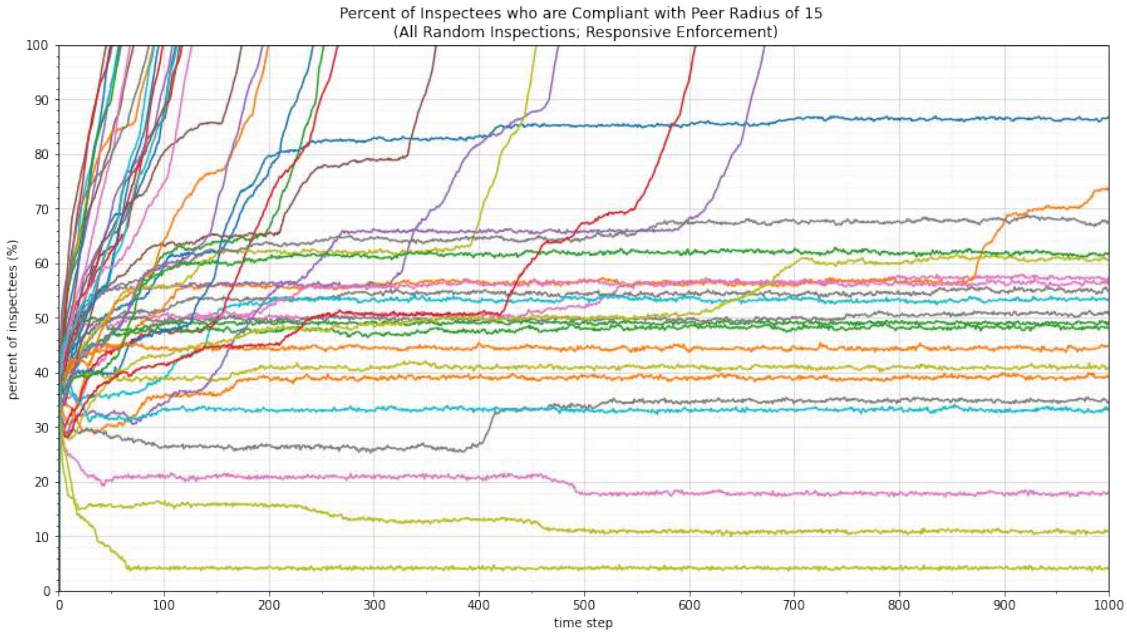

(a) Compliant inspectees.

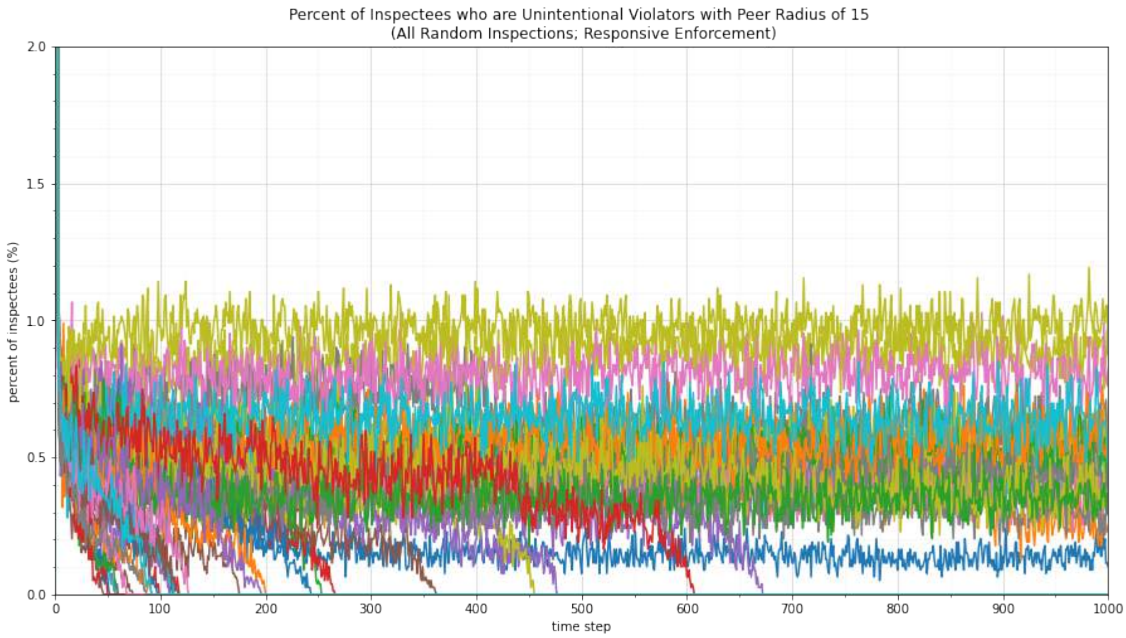

(b) Unintentional violators.

Figure A.72: Non-compliant inspectee breakdown with r_{peers}=15 (RE, AR).

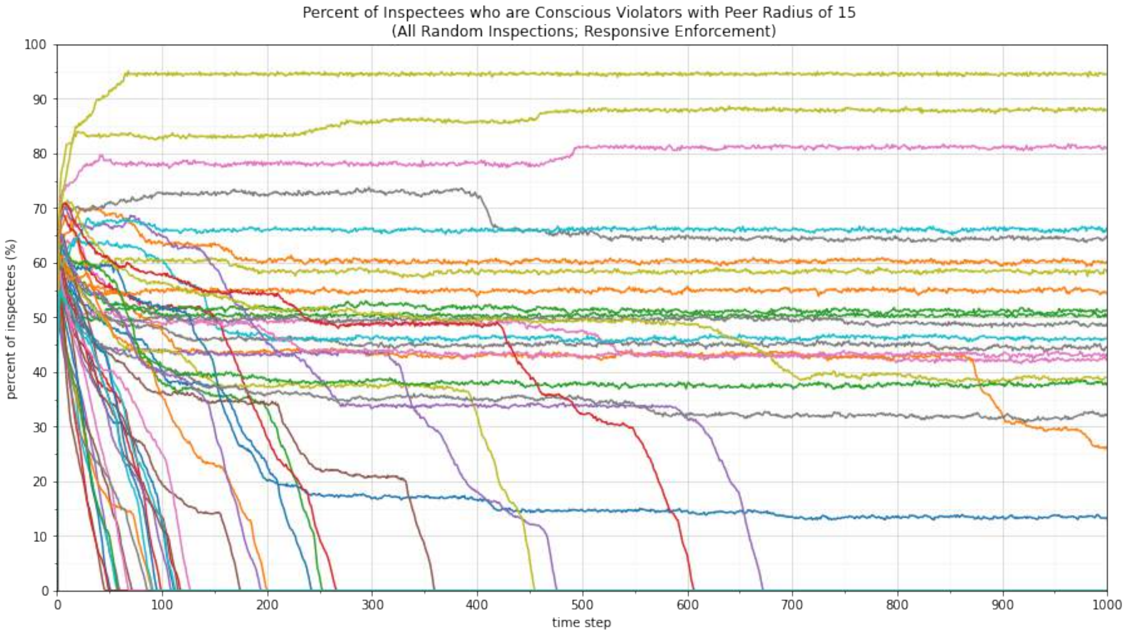

(c) Conscious violators.

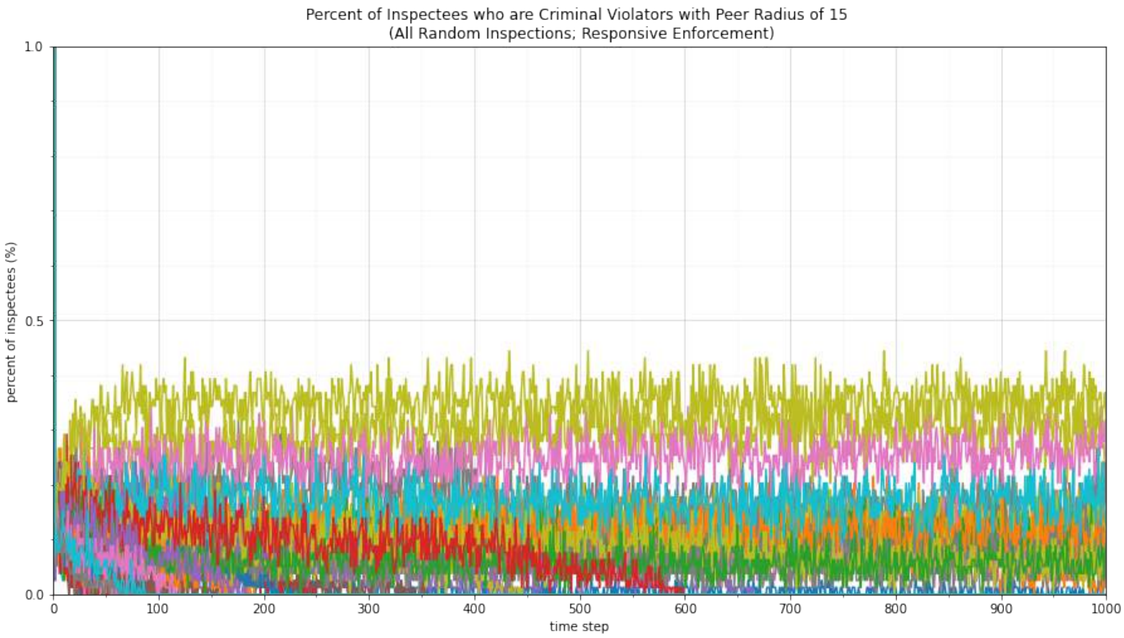

(d) Criminal violators.

Figure A.72: Non-compliant inspectee breakdown with r_{peers}=15 (RE, AR, continued).