A Risk Analysis and Data Driven Approach to Combating Sex Trafficking

by

Julia Coxen

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Industrial and Operations Engineering)
in the University of Michigan
2021

Doctoral Committee:

Professor Seth D. Guikema, Chair
Professor Bridgette Carr
Associate Professor Matthew F. Dabkowski, US Military Academy
Professor Nadine B. Sarter
Professor Lawrence M. Seiford

Julia O. Coxen

juliaoh@umich.edu

ORCID iD: 0000-0002-5314-8178

Dedication

This dissertation is dedicated to the two small humans who patiently waited for playtime or for a video to be restarted as their mother labored over this dissertation. I know that you are always watching and my hope is that my work inspires you to stay curious and value learning in and beyond the classroom. We are so proud of you both. ILY, AFNW

Acknowledgements

I would like to acknowledge the victims of sex trafficking and their continued suffering during and beyond the pandemic. May this work play even a small part in centering their voices and educating those in positions of influence to reverse the growth of this crime.

Acknowledgement is due in large part to my advisor, Dr. Seth Guikema, who graciously took a Ph.D. student from out of the blue under his wing. After many years in the Army I have been fortunate to meet incredible leaders who have indelibly changed the course of my life; I count him chief among them. If it were not for his compassionate leadership and steadfast guidance I would not be where I am today. I am intensely grateful for the way he not only gently pulls and pushes his students toward a goal, but walks alongside us and empowers us in this journey.

I would like to thank my committee members, Bridgette Carr, Colonel Matt Dabkowski, Ph.D., Dr. Nadine Sarter, and Dr. Larry Seiford for their precious wisdom and time. I am extremely grateful for their insights and mentorship. I learned a great deal from my collaborators Vanessa Castro and Glen Bredin and am grateful for their contributions. This work was also funded in part by the General Omar N. Bradley Research Fellowship in Mathematics.

Lastly, I need to acknowledge my two families. Thank you to the small family of dungeoneers who shuttled me, lifted me up, and most importantly taught me how to be a researcher—special thanks to Luke, Tessa, Erkin, Anna, Daniel and the lab. Thank you to my first family and heroes, Myoung Shook and Seung Suk, who made me believe anything is possible. Thank you to Ken and Sherry for your unflappable support and positive attitude. Thank you to Craig for his bottomless patience and devotion to our family—you are my favorite.

Table of Contents

List of Tables

List of Figures

List of Appendices

Abstract

Sex trafficking is a heinous criminal act that compels victims in the United States and worldwide to perform commercial sex acts through force, fraud, coercion, or age (TVPA, 2000). This dissertation takes a risk-analysis and data-driven approach to attain a better understanding of the problem, with the goal of showing that such an approach can help comprehend misallocation of resources, reform policy, reinforce social services, or support populations vulnerable to sex trafficking.

Sex trafficking is a complex problem and must be studied both qualitatively and quantitatively in order to provide those in a position of influence with an improved basis for decision-making. In Chapter 2 of this dissertation, I outline the risks associated with sex trafficking and suggest that risk analysis tools can be useful for antitrafficking efforts, as they can provide context-sensitive, empirical knowledge as well as a way to communicate neutrally about a charged topic.

Building on the understanding of this complex crime, in Chapter 3 I analyze online commercial sex work advertisements to draw conclusions about the COVID-19 pandemic's impact on sex trafficking, showing a measurable impact of the pandemic-related stay-at-home orders on advertising, and likely on the vulnerability of at-risk populations to trafficking. Finally, in Chapter 4 I use data collected by myself and a collaborator on sex work advertisements as a basis to explore three quantitative methods for detecting anomalies in time-series data. Based on the results of this sex trafficking case study, I evaluate the benefits and drawbacks of each method for risk-based decision-makers and discuss how these methods can be integrated into a broader risk framework.

This dissertation contributes to the field of sex trafficking research by offering improved methods for detecting anomalous behaviors in the system and advancing the application of these techniques for the risk analysis community. Although they are specifically designed for sex trafficking, analysts can apply these methods to many of the risk-related challenges of our future.

Chapter 1 Introduction

1.1 Sex Trafficking Overview

In the United States and abroad, human trafficking is a crime in which victims are made to work against their will. Being forced to perform sex work is only one form trafficking can take. Other forms of trafficking include debt bondage, domestic labor, forced marriage, and even state-imposed forced labor. In 2016 human trafficking in some form affected approximately 40.3 million victims (ILO, 2017). Approximately 25 million were in forced labor, including sexual exploitation (ILO, 2017). Although experts disagree on the magnitude of the problem due to the lack of homogeneous data, one thing is certain—the problem of sex trafficking is growing, and there is more occurring today than in past decades (Polaris, 2020; Savona & Stafanizzi, 2007).

The enormity and hidden nature of sex trafficking introduce a complex set of challenges for those attempting to combat its prevalence. Researchers may find it difficult to discover how demand is changing, where sex trafficking is growing fastest, and how laws intended to combat sex trafficking actually perform. A comprehensive bibliography of human trafficking research in 2015 determined that a substantial amount of qualitative research exists on human trafficking in social science, law, and medicine, yet there is a dearth of quantitative approaches to combating human trafficking. Much of the cited empirical research is quantitative, but only insofar as it pertains to evaluation, comparison, or description (Gozdiak et al., 2015). This characterization of the human trafficking research landscape is consistent with the research on sex trafficking. This lack of quantitative knowledge is due in some part to the clandestine nature of the crime and the muddled interpretation of trafficking definitions. These lead to a lack of empirical knowledge on which to base policy and enforcement efforts. Sex trafficking is often conflated with other related problems, such human smuggling and missing children (LCHT, 2017; Polaris, 2020). In reality, sex trafficking is a unique crime with its own patterns and challenges to prevention. Rigorous qualitative and quantitative analysis are needed to address the unique manifestations of sex trafficking and contribute to finding solutions.

This work addresses the gap in research by offering a risk analysis and data-driven approach to combating sex trafficking using publicly available information.

I organize the research according to three main objectives:

1. Understand the risks of sex trafficking: Show how a risk analysis tool can be used in the sex trafficking domain, and describe the characteristics a risk analysis tool should have in order to be used in the context of sex trafficking.
2. Investigate possible sex trafficking anomalies associated with a major social disruption and the factors that may contribute to potential increased trafficking.
3. Critically compare anomaly detection univariate time-series algorithms with synthetic data, generalize methods for risk analysts, and then apply them to a sex trafficking case study.

The main body of this dissertation is an ensemble of three papers, each of which aligns with one of these objectives. Combined, they form an interconnected body of work that contributes significantly to the risk analysis and sex trafficking literature. The remaining sections of this chapter will cover the motivations and objectives for this work.

1.2 Research Motivation

In a survey of law enforcement agents conducted by the U.S. Department of Justice, a positive correlation appears between agents' knowledge of human trafficking and perception of the crime's seriousness (Clawson, 2006). However, many lawmakers and law enforcement agents know little about this form of crime, or believe inaccurate myths about it (Clawson, 2006). My choice to take a data-driven and analytical approach in this work is motivated partially by my desire to provide an improved basis for knowledgeable decision-making, and to illuminate the anatomy of this complex issue without overstating the seriousness of the problem. These goals are especially important to me as a researcher in the United States. The United States plays a role in proliferating the nefarious practices of the multinational sex trafficking machine by providing demand for trafficked sex acts as well as supplying traffickers and victims of the crime. While researchers around the world continue to analyze effective strategies and governments struggle to legislate in a way that decreases the prevalence of sex trafficking, the United States has an opportunity to lead by example.

Policy-makers in a position to influence sex trafficking decisions need tools supported by data to make the best use of scarce resources like government funding for social services or policing. This applies not only to sex trafficking, but other complex risk-based problems. The ideal approach to all of them has the same basic outline. First, we need to understand the problem and identify the risks. Next, a risk-based decision maker needs to know when to be concerned: When do anomalies occur, and to which parts of the data do they pay special attention? They can reinforce this knowledge by examining the various political, socioeconomic, and environmental factors that influence the risk. Lastly, they should gather data about the effect of any attempted solution to determine empirically whether it has the desired effect. For any of these steps to be effective, researchers must understand the array of anomaly-detection methods that exist in their field and be able to place confidence in the methods they use to detect anomalies. This body of work contributes to all these aspects of risk detection; although not a formal decision-support tool, it provides a framework for the advancement of many risk-related challenges.

Motivation for this research is deeply personal. Sex trafficking and systemic racism against the Asian-American population have recently become part of public consciousness in light of the Atlanta massage parlor shootings on March 16, 2021. As an Asian-American, I feel deeply for those affected and believe this work is important in deconstructing the mechanisms that keep victims in sex trafficking situations or push vulnerable populations into sex trafficking. As an active-duty Army officer, I am also aware of the role my profession has played in the proliferation of sex trafficking (Hughes et al., 2007). Although academic research on the topic is sparse, many news articles describe the current normalization of commercial sex and sex trafficking as having roots in the military's overseas presence. Reflecting on my many deployments and interactions with the native and third-party national workforces common to many overseas bases, I wish I could have done more for those whom I naively assumed were not being exploited. It is nearly impossible to avoid patronizing all forms of human trafficking in our daily lives (Anthony, 2018), but this work will play a small part in correcting the quantitative literature gap and attempt to loosen the entanglement between sex trafficking and the military.

1.3 Project Objectives

This work can help policy makers develop mechanisms to decrease the incidence of sex trafficking and its severe harm to victims. We build on much of the work of the well-established risk analysis and industrial engineering disciplines to enhance our understanding of the risk-related challenges of sex trafficking and advance the detection of anomalies in context of sex trafficking as well. Table 1.1 summarizes each chapter's intellectual contributions and areas of focus.

Table 1.1: Intellectual Contribution and Dissertation Chapters' Methodological Focus

Chapters	Intellectual Contribution	Methodology focus
2	Risk analysis as a tool to combat sex trafficking	risk analysis foundations
2,3	Holistic evaluation of social disruption on a system based on obfuscated data	interdisciplinary data science
3,4	Close the gap in quantitative sex trafficking research	anomaly detection, algorithm comparison
4	Advancing univariate time-series anomaly detection	semi-supervised learning, change point detection, statistical process control

The sections below will briefly summarize the objectives for this ensemble of papers. The focus of this work is on the sex trafficking component of human trafficking; we are careful not to conflate any of its findings with labor trafficking, human smuggling, child pornography, or missing children. These are extremely important topics, but careless equivocations can have harmful effects on victims. We also avoid conflating sex trafficking with consensual sex work, while acknowledging that the two can be difficult to disentangle in practice since the trafficked status of a sex worker is frequently concealed from the public and potential clients (see Chapter 3).

1.3.1 Project 1: Risk Analysis as a Critical Tool to Combat Sex Trafficking

The objective of Chapter 2 is to show that risk analysis is a critical tool in combating sex trafficking. This research first aims to overview existing approaches to regulating sex work. This work goes on to categorize the risks of sex trafficking into risks to public health, security, and community in order to establish an understanding of the problem, before suggesting that a risk analysis tool must exhibit four key characteristics in order to be effective in the sex trafficking domain. The analysis makes comparisons with the well-documented domain of intelligent adversarial risk analysis and discusses lessons learned from terrorism risk analysis that can help policy makers unpack the complexities of sex trafficking. We are not aware of any other work that intersects with the field of risk analysis and sex trafficking; this work is thus unique in that sense.

1.3.2 Project 2: COVID-19 Pandemic's Impact on Online Sex Advertising and Sex Trafficking

The second project of this research dives deeper in to the risks created by sex trafficking demand in the United States as it pertains to major social disruptions such as the COVID-19 pandemic. Mainstream media often cite other major socially disruptive events such as the Super Bowl as a cause for spikes in demand and supply for commercial sex that drive up the prevalence of sex trafficking (Martin & Hill, 2019). These claims are rarely supported by evidence, and their sensationalism is a symptom of our failure to understand systematic drivers of sex trafficking. This project's objective is to examine the COVID-19 pandemic's effects on sex trafficking, and to quantify the impact of stay-at-home orders on the prevalence of both advertisements associated with independent commercial sex work and advertisements that we flagged as potentially related to sex trafficking. In this project I collaborate with Vanessa Castro, Bridgette Carr, Glen Bredin and Seth Guikema to determine research objectives and methods. We advance on existing methods for data handling and data interpretation to derive value from data that have been purposely obfuscated. Our research takes steps towards closing the gap of quantitative literature in the sex trafficking domain. To my knowledge, all previous work that measured sex trafficking prevalence in this way was conducted prior to the passage of U.S. SESTA-FOSTA legislation that significantly affected online commercial sex work (Chamberlain, 2018). This research may be the first on the subject since that landmark legislation, and also takes steps towards closing the gap of quantitative literature in the sex trafficking domain.

1.3.3 Project 3: Time-Series Anomaly Detection Algorithm Comparison for Risk-based Decision-Making

Risk analysts can benefit from further exploration of data sets similar to the one explored in Chapter 3. The research in Chapter 4 evaluates potential methods to glean useful information from such data by multiple algorithms for detecting anomalies (unusual events) in time-series data. The risk-related challenges of our future range from climate change to critical infrastructure systems (Greenberg et al., 2020). However, such techniques have not previously been applied to the problem of sex trafficking. The growing volume of temporal data streaming in from sources such as sensors and wearables makes it even more important for risk researchers to be prepared to recognize anomalies and where to draw their attention to risk.

The objective of Project 3 is thus to critically compare univariate time-series anomaly detection techniques, focusing on the detection power and robustness of machine learning, change point detection, and statistical quality control. We hypothesize that different algorithms will be best suited to detect different types of anomalies. We conducted experimental tests of these three approaches using a synthetic data set representative of the count data from the online advertisements from Chapter 3 of this dissertation and apply them to our real-world data set, ultimately concluding that the change point algorithm detects risks most effectively in this context. The literature is rich with algorithm comparisons, but Project 3 (Chapter 4) expands the anomaly detection literature by using concrete examples crossing over multiple disciplines, by contextualizing these techniques for the risk analysis community, and by offering a case study with real-world data on suspected indicators of sex trafficking.

Chapter 2

Risk Analysis as a Critical Tool to Combat Sex Trafficking

Sex trafficking is a global crisis that continues to be difficult to detect and prosecute. This chapter will examine current policies in place globally to combat sex trafficking, and argue that risk analysis, while currently underused, is a critical tool in combating it. The crime's intractable nature often marginalizes the problem and makes holistic evaluations challenging. Understanding the risks will help establish common ground and help policy-makers develop mechanisms to decrease the incidence of the crime and its effect on victims. Finally, we outline the features a risk analysis tool must have in order to make a significant impact in the sex trafficking domain. A risk analysis tool would need to account not only for the adaptive nature of its adversary, but also for uncertainty around the problem. The tool must address the continuous nature of the risk and the presence of multiple players. Although further research is needed to perfect them, risk analysis methods can contribute significantly to stopping this problem.

Keywords: sex trafficking, risk analysis, risk modeling, human trafficking

Note: This research will be submitted to the *Journal of Human Trafficking*.

Co-authors include Bridgette Carr and Seth Guikema.

2.1 Introduction

Sex trafficking is a global crisis that that is exceedingly difficult to detect and prosecute. It is difficult to quantify the magnitude of the problem because of the clandestine nature of the crime (Konrad et al., 2017). Prosecution of a sex trafficker can be challenging not only because their crimes are difficult to detect, but because prosecution often relies on the cooperation of the victim, who may fear retaliation (Riegler, 2007). The existing paradigm for fighting sex trafficking is deeply grounded in a law enforcement approach that focuses on detecting and prosecuting commercial sex workers (McCarthy, 2012). We argue in this paper that risk analysis can offer a critical supplement to this way of thinking, broadening the perspective and allowing policy-makers and law enforcement to take a risk-based approach to combating sex trafficking. However, to achieve this, additional research is needed on how to detect, mitigate, and prevent the risks of sex trafficking.

The prevalence of sex trafficking crimes is staggering; the number of victims across the globe is estimated at 15.4 million (ILO, 2017). These include many who do not fit the stereotyped view of a trafficked sex worker as a victim of kidnapping or smuggling across long distances. While the word “trafficking” implies transportation, no movement is necessary: Sex trafficking includes any situation in which a victim is forced to perform commercial sex work through force, fraud, coercion, or age (Trafficking Victims Protection Act of 2000). For instance, an independent commercial sex worker can be blackmailed by a pimp who threatens to reveal their criminal line of work unless they do as required; this can happen in a victim’s own home, neighborhood, or city.

Sex trafficking is widely feared and viewed as a threat to community safety. Nevertheless, there is little in the way of research literature that empirically assesses the signs of risks and most effective interventions. Despite motivation, researchers may still find it difficult to conduct work in this domain because of the lack of reliable data. Experts struggle to provide basic prevalence estimates and measure the scope of the problem (Gozdiak, 2011). Nevertheless, risk analysis techniques may be useful even in the absence of such information through the careful incorporation of expert opinion. For example, one of the advantages of using risk analysis as a tool to combat sex trafficking is its well-established ability to provide risk estimates on the basis of expert judgment without extensive data sets (e.g., Apostolakis, 1990). Ending sex trafficking will take a broad approach integrating academic and societal contributions (such as education and social services) as well as extensive agency coordination to curate accurate statistics for risk modelers. The field requires academically rigorous and intellectually robust examination of sex trafficking from a risk analysis perspective in order to provide sound recommendations to decision-makers.

The risk management approach is distinct from other common approaches used to combat social ills. In this paper, we provide a high-level overview of some current policy approaches used to combat sex trafficking. We then outline the requirements risk analysis tools would need in order to be effective in the sex trafficking domain. We then close with suggestions for moving risk analysis forward to help in the fight against sex trafficking.

In the progress of writing this paper, we interviewed law enforcement experts, current and former sex workers, and antitrafficking experts to gather general background information. We are not attributing any feedback to a particular interviewee in order to honor their need to remain anonymous.

2.2 Overview of Policies to Combat Sex Trafficking

The existing policy setting is largely focused on preventing both sex trafficking and independent commercial sex work, with some exceptions. Countries have legalized commercial sex work to differing degrees, with different approaches to legalization, and varying degree of focus on sex workers versus sex buyers. Laws do not always differentiate clearly between sex trafficking and independent commercial sex work. A population of independent sex workers exist who are not trafficked; recall that sex trafficking is distinguished from sex work by the inclusion of force, fraud, or coercion for adults or if the sex acts involve a minor. The policies intended to curb one form of sex work can significantly affect the other. For instance, shutting down a website that facilitates sex trafficking can create difficulties for independent commercial sex workers who advertise their services on those same websites.

This section focuses on how commercial sex work is regulated and is not comprehensive of all the existing methods to decrease the prevalence of trafficking. For example, community-level efforts to decrease demand, efforts to reduce sex tourism, car confiscation programs (in which law enforcement officers confiscate cars belonging to criminal sex buyers), and faith-based programs are not addressed here (Hughes, 2004). While this section focuses on variation, the current status quo in the vast majority of countries reflects a fully criminalized model where a ban exists on the buying and selling of sexual services (McCarthy, 2012). Legal frameworks in which sex work is legal for some or all parties are variations from a prevailing model in which both sex work and sex trafficking are punishable offenses.

2.2.1 Legalization of Commercial Sex Work

One regulatory approach is to legalize commercial sex work. There exists, however, disagreement in the literature and among policy-makers about the effects of legalization on the prevalence of sex trafficking. Legalization has the potential to decrease the incidence of sex trafficking and risks to victims by placing the activity within the boundaries of regulatory oversight, leading to a substitution of legal, regulated commercial sex for commercial sex provided by trafficking victims (Cho et al., 2013). However, legalization could also normalize sex work and increase demand for commercial sex beyond what legal sex workers could provide, increasing sex trafficking (Cho et al., 2013). Other literature suggests that coupling regulatory agencies with law enforcement in the fight against trafficking does not increase the transparency necessary to detect or decrease sex trafficking (Huisman & Kleemans, 2014). One global analysis of over 150 countries (not including the United States) suggests that the scale effect eclipses the substitution effect and that sex trafficking increases in those countries that legalize sex work (Cho et al., 2013), but more work is needed to better quantify the impacts of legalization.

2.2.2 Decriminalization

In a legalized setting, sex work is lawful under certain conditions. For example, in the U.S. state of Nevada, street sex work is not permitted, whereas sex work in brothels is legal, but must adhere to state regulations and licensing protocols (Weitzer, 2015). Legalization focuses on the regulatory component of the commercial sex enterprise and establishes specific circumstances under which it can be pursued, just as restaurants can operate only if they follow health codes and other rules.

Decriminalization is distinct from legalization and focuses on the prosecutorial nature of the crime. Decriminalization removes the punishment associated with selling sex services, buying sex services, or being involved in the third-party management of the services. It is a shift in policy that moves the focus from punitive laws that punish all facets of sex work to policy and strategies that focus on protection of sex workers (Albright & D’Adamo, 2017). The removal of punishment for those providing sex services makes it possible for sex workers to report abuse and exploitation. This effect is particularly amplified in sex trafficking victims of color whose access to justice is already compromised due to systemic racism (Chu & Glass, 2013). These vulnerabilities are notable in the context of sex trafficking, where victims exploited through force, fraud, or coercion are reluctant to report exploitation or are unaware that they are being abused (Kara, 2017).

Prostitution is legal or decriminalized in many countries around the world, including Canada, Mexico, Australia, New Zealand, France, Germany, and India, although often with strict laws around where and how sex work can be engaged in and advertised. Critics of decriminalization and legalization claim that the removal of laws related to sex work normalizes exploitation, increases demand, and places sex trafficking victims at greater risk (McCarthy, 2012). On the other hand, decriminalization advocates argue that the criminalization of sex work increases violence, erodes trust, increases vulnerability, and reinforces a stigma that continues to marginalize sex workers (Albright & D'Adamo, 2017). Regardless of where they stand on this debate, policy-makers should insist on policies that prevent sex trafficking and support victims of physical abuse in order to protect sex trafficking victims, while also considering how criminal changes can affect sex workers' rights and the possibility of exploitation, marginalization, and poverty.

2.2.3 Partial Decriminalization Method (Nordic Model)

An alternate model attempts to protect sex workers by criminalizing buying, but not selling, of sexual services. In 1999, Sweden passed legislation that criminalized sexual services only for the buyer, based on the premise that without demand there is no supply. Previously in Sweden, both buying and selling sex had been decriminalized and were not subject to any law. Given that the demand for commercial sex does not usually discriminate between voluntary and trafficked workers (as clients often do not know which is which), a policy that addresses the demand side of the enterprise affects both commercial sex and sex trafficking. The Swedish model turned into the Nordic model in 2009 when Norway adopted a similar policy outlawing the purchase of sexual services. Finland followed suit 5 years later, building on their previous 2003 legislation that outlawed the buying of sex in public spaces (Hughes, 2004).

In general, the Nordic Model makes self-employed individual sex work legal while also making a direct link between sex work and sex trafficking by affecting them both to a similar degree (Ekberg, 2004). This approach has attracted some support. Some experts believe that the prosecution of clients makes the purchase of sexual services less appealing, decreasing demand and thus making it less profitable for traffickers to establish networks in these areas. The National Criminal Investigation Department of Sweden (Raymond, 2004) reported that the decriminalization of sex work with the simultaneous criminalization of sex buying has prevented traffickers from making money as quickly as they could in other countries where sex work is more broadly criminalized. This was likely because the buyer's invisibility had been compromised, forcing traffickers to use multiple locations (e.g. hotels and apartments) to provide services. In this one case, the increase in cost has been found to be enough of a deterrent to reduce the overall incidence of sex trafficking (Raymond, 2004).

While the Nordic model has attracted support as a seeming compromise that reduces demand for sex work without punishing sex workers, some experts have cautioned that it is ineffective or even counterproductive in preventing trafficking. Notably, clients may be "deterred from reporting potential evidence of sex trafficking and exploitation to the police, because they themselves fear criminal prosecution," and victims may be "more reluctant to report their exploitation to the police for fear of reprisals from traffickers, and because of their inherent distrust of the authorities" (Kingston & Thomas, 2019). The continued policing of the sex trade (even if only to arrest clients) may drive sex workers to more isolated, less safe locations and result in them having less negotiating power to insist on safe sex practices (Chu & Glass, 2013). Other researchers have argued that in practice, sex workers are not protected by the legal status of their work: "The regulation of commercial sex has shifted from prostitution to immigration policies….Nationals are provided social welfare policies to assist exit from commercial sex such as therapeutic counseling, whereas foreigners are excluded from state services and targeted with punitive measures, like deportations and evictions" (Vuolajarvi, 2019). It is thus argued by some that complete decriminalization is more effective at ensuring legal protection for trafficking victims.

2.2.4 Demand Reduction

In the United States, sex work is legalized in some form only in Nevada, and has not been formally decriminalized in any state or municipality. Short of adopting a decriminalization policy, some U.S. cities chose to incorporate the principles of other decriminalization efforts by performing demand-side interventions (DOJ, 2012). As with the Nordic model, governing officials see this "buyer beware" approach as a way to reduce demand while avoiding the prosecution of victims of sexual exploitation (Kroman, 2018). This strategy differs from partial decriminalization in that no legal conditions change, but local governments and law enforcement choose to institute tougher policing towards buyers and traffickers. This strategy has the potential to reduce demand and induce more sex trafficking victims to report abuses, but provides no legal protections if victims come forward as having provided sex services.

One example is the U.S. city of Seattle, where the selling and buying of sex are still both illegal, but law enforcement strategy has been far more aggressive about prosecuting buyers, focusing on an "end demand" approach (Kroman, 2018). The *Seattle Times* reported some promising early successes, with almost three times as many buyers as sex workers being arrested than across the county in the program's first year (Welsh-Huggins, 2017). Across the Unites States and Canada, some cities are simultaneously implementing "john schools" to educate arrested buyers through required programs that focus on the impact of sexual exploitation and trafficking on victims (Hughes, 2004). This locally enforced strategy does not directly fight sex trafficking or provide voluntary or coerced sex workers with a legal remedy against persecution, and still results in the arrest of sex workers.

An exception to the United States' criminalization of sex work existed in Rhode Island for a 30-year period beginning in 1980; in this time period the state law forbid the buying and selling of sex only outdoors. This experiment in decriminalization happened unintentionally, as lawmakers, facing an increase in sex work in the state, moved to make prostitution a misdemeanor instead of a felony in hopes of speeding up convictions (Gordon, 2017). In rewriting the law, they identified only street sex work as a crime, allowing online escorting and other indoor sex work to take place legally. This loophole was closed in 2009. Rhode Island's recriminalization of sex work has been criticized by some; an ACLU report notes that "during the period of inadvertent decriminalization of 'indoor prostitution' in Rhode Island from 1980 through 2009, there was a 30 percent decrease in reported rape offenses against sex workers post-decriminalization" (ACLU, 2020, p. 6). Similar decriminalization laws have been proposed in states such as New York, Vermont and Washington, D.C., but none have yet passed. Another form of informal decriminalization occurs when prosecutors decide not to prosecute sex work offenses; for example, Baltimore City prosecutor Marilyn Mosby announced in March 2021 that the Maryland city would stop prosecuting minor drug possession charges, sex work offenses and other nonviolent offenses in an effort to focus on violent crime (Battaglia, 2021).

2.3 Risks of Sex Trafficking

Before discussing how risk analysis can play an important role in the fight against sex trafficking, we give a brief overview of the risks that sex trafficking can pose to: (1) public health, (2) security, and (3) communities. In addition to directly harming victims, sex trafficking is associated with other crimes including gang crime, drug and property crimes, and organized criminal operations in many places such as southeast Asia, Europe, and the United States (Shelley, 2012). Some of the connections we make to the risks of sex trafficking have roots in these adjacent crimes as well as the complex nature of the sex trafficking enterprise.

2.3.1 Risk to Public Health

Sex trafficking poses significant health and safety risks to its victims. They may experience sexual, physical, and emotional abuse by both traffickers and buyers, while being exposed to other workplace hazards such as risk of physical harm, document confiscation, and poor working conditions (e.g., denied access to contraception, lack of safe food, or substandard living quarters). A study of sex trafficking victims in the European Union found that 76% reported physical violence, 90% reported sexual abuse, 76% reported restrictions on freedom, and 57% reported signs of post-traumatic stress (Zimmerman et al., 2008). Zimmerman (2011) lays out the range of abuses to which sex trafficking victims are vulnerable, including poor working conditions, poor living conditions, physical abuse, sexual abuse, psychological abuse, and restricted movement. Interpersonal violence is a common form of coercion by the traffickers, and victims may experience violence at the hands of buyers as well (Zimmerman et al., 2008). Victims are also vulnerable to harm during high-risk transportation such as desert or ocean crossings, and to violent retribution of traffickers if they attempt to escape or are forced to return to trafficked sex work.

In addition to working and living in poor conditions and being forced to have sex, sex trafficked victims are often denied the use of condoms, increasing the spread of sexually transmitted infections and diseases. Access to medication and healthcare is limited for victims; therefore, they will go untreated and continue to pose a risk on the local population and to any future partners. A 2009 study of repatriated sex trafficking survivors in Nepal found a higher incidence of tuberculosis among victims, and nearly 9 in 10 individuals who developed TB were co-infected with HIV (Dharmadhikari et al., 2009). A study of 107 U.S. sex trafficking survivors ranging in age from 14 to 60 reported that 67.3% of victims experienced some type of sexually transmitted disease, such as chlamydia, gonorrhea, or Hepatitis C (Lederer & Wetzel, 2014). The spread of disease has been most recently illustrated by the COVID-19 pandemic. As social distancing measures were put in place across the United States, news media reported a decline of in-person sex work (Breslin, 2020). Although no data exist yet on the rate of SAR-CoV-2 virus spread through sex work or sex trafficking, initial anecdotal reporting suggests that the fear of contracting the virus from a sex worker is strong enough to affect sex work demand (Breslin, 2020). Unwanted pregnancy is another outcome of unprotected sex. Despite complications of reporting, the Lederer study reported that “pregnancy, miscarriage and abortion were common experiences for survivors in the study” (Lederer & Wetzel, 2014). There exists no doubt that the harm inflicted on sex trafficking victims is severe.

2.3.2 Risk to Security

Sex trafficking not only imposes violence on victims but poses a risk to national and local security. It is a criminal enterprise driven by "immense profitability with minimal risk" (Kara, 2017), and this enterprise is often associated with additional criminal activity. Much of the profit from trafficking is routed into criminal organizations, further harming the safety and security of communities (Kara, 2017). Sex trafficking strengthens organized crime, which puts overall security at risk by sustaining criminal activities such as "drug trafficking, human smuggling, money laundering, and document forgery" (Jones et al., 2007). Additionally, a recent report by the Council of Foreign Relations finds that "human trafficking can fuel conflict by enabling armed and extremist groups to raise revenue and expand their power and military capabilities" (Bigio & Vogelstein, 2019a). Bigio and Vogelstein give examples such as Boko Haram, the Islamic State, al-Qaeda, Al-Shabab, and other organizations that use sexual violence to terrorize their victims into compliance and to displace civilians from strategic areas.

Although not all traffickers operate as part of an organized criminal network, many work in tandem with criminal cartels and terrorist groups to further the reach of both organizations (Howard, 2013). In his article about the nexus between trafficking and terrorism, Howard discusses the many similarities between international criminal and terrorist organizations. They both use violence, fear, and coercion in areas where government control is weak and borders are porous. Violent extremist organizations such as the Islamic State in Iraq and Syria (ISIS) and Boko Haram use sex trafficking to generate revenue (Bigio & Vogelstein, 2019a). The United Nations Security Council has considered the relationship between sex trafficking and terrorism (United Nations, 2019), finding that terrorists use sex trafficking victims to attract and retain fighters. Their studies also show that sex trafficking is used as a terrorist tactic to intimidate populations and decimate communities, to institutionalize sexual violence, and as a driver for recruitment. It is dangerous to conflate trafficking and terrorism, since it could dilute the perceived importance of the stand-alone issue of sex trafficking. Additionally, sex trafficking within U.S. borders is rarely connected to organized criminal activity. However, there is an undeniable connection between terrorism and trafficking that threatens our security in an international setting.

2.3.3 Risk to Community

The presence of sex trafficking can cause things to go wrong in a community, undermining society and quality of life. These effects can take myriad forms, but three are especially salient. First is the proliferation of fear, which can lead to poor, fear-driven decision-making by community members and policy-makers. Second is the intersection of sex trafficking with legitimate businesses that could unknowingly support sex trafficking such as hotels and financial services. Lastly, traffickers create conditions that place more strain on essential social services. With these factors in addition to the discussed risks to public health, communities are clearly worse off with the presence of sex trafficking, though more research is needed to better understand the risks.

As is the case with many risks, it can be very challenging for experts to communicate sex trafficking risks to the public and policy-makers, and the social amplification of risk can have significant impacts on policy and decision-making (Kasperson et al., 1988). As a direct result of the complex nature of sex trafficking, the absence of knowledge about this clandestine crime often leads to hyperbolic levels of dread and fear or apathy toward the crime. A study of print and broadcast media on the subject of sex trafficking found that news was “largely episodic, focused on crime and policy frames, privileged official sources, and neglected survivors’ voices” (Johnston et al., 2015). We can infer that an infrequent news consumer or even an ordinary middle-class citizen might not be well-informed about sex trafficking or see it as a serious problem, relegating it to the same category as other existential threats such as terror attacks that may never affect them. On the other hand, an individual who consistently consumes media may inflate the danger of sex trafficking and view it as a threat to their personal safety, like other random crimes such as muggings. Flawed understandings of how sex trafficking operates and who is at risk may dilute the understanding of it as a human rights and public health issue (Johnston et al., 2015). Provocative stories of kidnapping and forced prostitution litter the news media, which reinforces the misperception that sex trafficking victims are predominately young white girls snatched from their neighborhoods and moved across borders (Johnston et al., 2015). This type of characterization not only is misrepresentative of sex trafficking and its victims, but also creates barriers to understanding. The fear proliferated in communities can lead to disproportionate policy responses or strategies that focus on educating the wrong population. Risk analysis can play a critical role in combating this with holistic risk-aware responses to the problem.

Sex trafficking does not only operate in the shadows, but involves many legitimate businesses and organizations. A recent Polaris Project report details the on-ramps to sex trafficking and their intersection with many industries including social media, financial services, hotels, transportation, healthcare, and housing/services for the homeless (Anthony, 2018), making these industries complicit (sometimes unknowingly) in trafficking. Although the report does not make scientific conclusions, the author emphasizes reports that "traffickers use banks to store their earnings, and buses to move their victims around; hotel rooms are integral to the operations of some sex traffickers, social media is a vital recruitment trawling ground for others" (Anthony, 2018). Not only do the criminal activities of sex trafficking make legitimate industries more vulnerable by association, they use established businesses as platforms to carry out their business plans (Anthony, 2018). Members of a community may thus be unknowingly patronizing establishments that support sex trafficking, magnifying the harm caused by the previously mentioned risks to public health and risk to security.

Lastly, sex traffickers create conditions in which victims need essential services in order to reintegrate back into society. As discussed previously, the harm they inflict on victims is severe. The few victims who are able to escape their exploitation should be offered essential community resources and social services to rehabilitate back into the community. These individuals are often in need of intensive physical and mental medical care, emergency shelter, and other resources such as job training to re-enter the community. Traffickers often use drugs and alcohol to coerce their victims and keep them enslaved (Shelley, 2012). This means that traffickers impose risk on the community by increasing the burden on drug/alcohol rehabilitation programs, shelters, humanitarian organizations, social workers, medical facilities, employment assistance, and more. Taken together, these risks to communities are difficult to quantify, but clearly are significant.

2.4 Risk Analysis is a Critical Tool in the Fight Against Sex Trafficking

Risk analysis, with its focus on identifying and intervening in the most significant sources of harm, can aid the fight against sex trafficking by providing an improved basis for policy and law enforcement decision-making. The first essential is an accurate perception of the magnitude of the problem. "Whether human trafficking is viewed as a serious problem or considered a priority crime can affect law enforcement's capacity to respond" (Clawson, 2006). A U.S. Department of Justice report on law enforcement responses to human trafficking included surveys of local, state, and federal law enforcement agents (Clawson, 2006). When asked about human trafficking as a priority, 58% considered human trafficking a high to very high priority, while almost half—20%

and 22%, respectively—reported that it was not a priority or "somewhat" of a priority. The authors suggest this assignment of low priority reflects low knowledge of the problem: "There was…a positive correlation between knowledge of human trafficking and perceived seriousness of the problem, suggesting more knowledgeable respondents are more likely to view human trafficking as a serious problem" (Clawson, 2006). Although this study is from 2006, no recent research exists to suggest the situation has changed. While this highlights the need for law enforcement education and training, it also highlights the need for deep understanding of sex trafficking and the proliferation of rigorous research on the topic.

But it is not enough for policy-makers to make trafficking a priority, and it may even do harm if their action is based on inaccurate beliefs. In order to curb the prevalence of sex trafficking, governments need deep knowledge of the problem and capacity to address the root causes and effects. Researchers and policy-makers should use the well-established and highly effective tools offered by the field of risk analysis to reduce the sex trafficking problem. In order for this to be effective, however, we must take a thoughtful approach to determine which features of risk analysis models we can apply to combating sex trafficking.

The exclusion of holistic research that systematically derives long-term risks from interventions can lead to unintended consequences. This is perhaps best represented in the sex trafficking domain by the passage of the Fight Online Sex Trafficking Act – Stop Enabling Sex Trafficking Act (FOSTA-SESTA), the U.S. House and Senate bill that passed in 2018 (Allow, 2018). The U.S. Congress passed a law targeting one particularly visible tool used by traffickers—Backpage.com—without considering the broader effects this legislation would put in motion. Many, including some law enforcement agents interviewed for this dissertation, argue that this law has pushed trafficking further underground, making it much harder to track and prosecute while creating significantly higher risks to both victims and independent commercial sex providers by reducing their ability to screen buyers (see also Chamberlain, 2018). A paper in Fordham Law Review reported additional unintended consequences that immediately followed the enactment: Sex workers were reported missing, and others faced tremendous increase in unwanted demands for trafficking from pimps (Chamberlain, 2018). A risk-based approach could have attempted to eliminate the nefarious sex trafficking activities from the websites while preserving the safety tools they provided for independent commercial sex workers.

Risk analysis methods can offer both an important perspective and key tools to improve policy-making and law enforcement interventions. The risk analysis process involves a holistic evaluation of the effects of intervention decisions. It acknowledges and explicitly addresses uncertainty about the future effects of actions, and offers many tools to help decision-makers allocate scarce resources to areas of highest priority. It can provide a better basis for understanding and communicating about the true risks of sex trafficking by providing fact-based, neutral communication about the scope of the problem, allowing for grounded discussions of policy and law-enforcement interventions.

However, for risk analysis to meet its potential to help in the fight against sex trafficking, we must carefully consider how well its current methods and knowledge match what is needed for use in the sex trafficking domain. That is, what must a risk analysis model be able to do? Many fields that make use of risk analysis are dramatically different from sex trafficking prevention and do not offer methods that can readily be transferred to this domain (for instance, risks to public safety caused by human error or extreme weather rather than criminal activity). Below, we use tools from the field of terrorism risk analysis as a basis of comparison because of (1) some similarities with the sex trafficking problem and (2) the significant established body of research that has been done in this area. However, there are significant differences as well. These model features build on the terrorism risk analysis foundation and can be applied to the sex trafficking domain, but more research is needed.

2.4.1 Necessary Features of Risk Models for Use in Sex Trafficking

2.4.1.1 The tool must address an adaptive adversary

The first and perhaps most obvious feature needed in a risk analysis model to combat sex trafficking is an ability to address adaptive adversaries. Most traffickers are driven by desire for monetary gain and will adapt their strategies in order to avoid detection and prosecution (Bouché & Shady, 2017). Risk models that do not account for their adaptive strategies will offer a misleading assessment of sex trafficking. Adversarial risk analysis methods have undergone significant development in the past 2 decades, largely in response to terrorism threats (Rios & Insua, 2012). While terrorism and trafficking are not interchangeable, they both involve an intelligent adversary seeking to do harm; as the literature on terrorism risk analysis has shown, it is critical to explicitly account for and model intelligent adversaries if we are to provide adequate support for resource allocation decision-making (Cox, 2009).

Consider as an example what might seem at first to be a “simple” problem: raiding and shutting down illicit massage parlors. It is relatively simple to determine which massage parlors offer illicit sexual services for sale based on public information (Feeney, 2013). However, according to law enforcement experts we interviewed, traffickers reopen illicit massage parlors quickly after a raid, sometimes moving them to the next county, and other times reopening in the same place with a new name. Similarly, when a major web page used for advertising commercial sex is shut down (i.e., Backpage.com), sex traffickers quickly adapt, switching to websites outside of U.S. jurisdiction (i.e., Rubmaps.ch). Shutting down these businesses may thus be a poor use of resources unless authorities have a plan in place to prevent them from reopening or establishing other platforms to advertise. Present law enforcement tactics in widespread use often fail to take into account this adaptive quality of traffickers and spend significant resources to shut down illicit businesses which may only be effective for a short amount of time. Traffickers are highly adaptive and switch tactics and locations in response to policy and law enforcement actions in order to maximize their profit. Any risk analysis tool meant to help in the fight against trafficking must be able to account for this adaptive nature of the adversaries in a way that provides direct support for decision making. Considerable research exists within the terrorism risk analysis literature on how to incorporate adaptive adversaries into risk analysis (Cox, 2009; Ezell et al., 2010; Guikema, 2012). This work can provide a starting point as the risk analysis community builds better modeling tools and frameworks to help combat sex trafficking.

2.4.1.2 The tool must address uncertainty

Addressing sex trafficking requires dealing with a great deal of uncertainty. Uncertainty can arise from the unpredictability of trafficker actions, lack of information on where, when, and how trafficking is occurring, and lack of knowledge about the purchasers of commercial sex. Much of this uncertainty is due to a fundamental lack of data. For instance, because traffickers operate illicitly, researchers and law enforcement do not know for certain who they are or how many people in a given area are actively involved in trafficking. Collecting the information needed to reduce this uncertainty is prohibitively difficult in many situations, though researchers have made progress in removing some uncertainty in what factors make populations more vulnerable to becoming victims of sex trafficking (Bales, 2007; Schwarz et al., 2020). We recognize that uncertainty also arises from the variability between communities, the variability in the players, and in the environmental, legal, and social structures where the sex trafficking resides.

Risk analysis tools used to support decision-making in this realm must represent both uncertainty and the degree of knowledge underlying the risk assessments in a way that decision-makers will understand. While probabilities can be used to represent uncertainties in many situations, they alone will not provide sufficient information to decision-makers. Advances have been made in representing the degree of knowledge underlying risk assessments (e.g., Flage et al., 2014), but additional research is needed to better understand how practical users of the risk assessments understand these knowledge assessments. This is particularly critical in the area of sex trafficking, as policy-makers and those in law enforcement are not generally accustomed to thinking about their problem in a risk-based way. A risk analysis tool must also be able to incorporate the descriptive and empirical work common in this field. One potential way to address this would be to use an agent-based model to represent the different individuals involved (e.g., traffickers, buyers, victims, and law enforcement; Huddleston et al., 2008, Paté-Cornell & Guikema, 2002). An agent-based model is a way of generating predictions when one is unable to make empirical observations. Groff et al. (2019) describe it as a scientific tool in simulation modeling that allows for the representation of policy and decision effects through the use of agents that simulate the real world. For instance, in modeling gang activity, analysts might model gang members' social network connections, while for street crime they might model a physical environment where criminals, victims, and law enforcement might interact (Groff et al., 2019).

Another potentially useful tool from risk analysis is the Bayesian network (Kiss et al., 2020). A Bayesian network is a probabilistic graphical model that represents the relationship among a set of variables that are related in some way, indicating which variables are random and which have a likely causative relationship. For instance, a Bayesian network might portray the relationship between diseases and symptoms; certain symptoms might have a known percentage of indicating cancer, and that probability might be higher if the patient is a smoker. Given information on the patient's symptoms and lifestyle factors, a Bayesian network could calculate the probability of various diagnoses. This tool allows expert knowledge and multivariate data to be combined in a visual way that is more accessible to decision-makers through the graphical network model.

2.4.1.3 The tool must address the continuous or nearly continuous nature of the risk

Much of the adversarial risk analysis research has focused on defending against one or a small number of discrete attacks against a well-defined, contained system (Gil et al. 2016; Cox, 2009; Rios & Insua, 2012; Merrick & Parnell, 2011). For instance, if analysts are trying to prevent airplane hijackings or airport attacks, the number of potential attackers is limited, and while even one successful attack would be catastrophic, there are not always active attempts being made to perpetrate such an attack. This is where a key difference exists between sex trafficking and threats such as terrorism. Sex trafficking is a persistent threat. Exploitation of victims, and thus trafficker adaptation, occurs continuously, not in discrete events as would happen with a terrorist attack (or many cybersecurity attacks), as traffickers rely on a continuous business presence to maximize income (Kara, 2017). Their activities never fall to zero, but rise and fall based on large-scale economic, governmental, law enforcement, and other activity. This means that the risks to community, health, and security are continuously varying in time. A risk analysis tool must be able to account for the continuous or nearly continuous nature of sex trafficking risk, and research advances are needed here.

Existing adversarial risk analysis models are not designed to address this type of continuous threat. For example, most game theoretic models of terrorism risk analysis and cybersecurity model single attacks (Paté-Cornell & Guikema, 2002; Cox, 2009, Merrick & Parnell, 2011) or a small set of discrete attacks (e.g., Levitin and Hausken, 2010). The modeling problem addressed in this research is substantially different from one in which all players can continuously change their strategies, sometimes in secret and often with different information sets. How do we model equilibrium behavior in such a case? How is learning or inferring from partially hidden information included? Is an equilibrium-seeking modeling approach the right framework for this problem, or is this better suited for an agent-based modeling type of approach? Research is needed to help answer these questions and to develop methods appropriate for continually varying threats in which the different players operate with different information sets.

2.4.1.4 Must account for multiple players

Finally, sex trafficking risk analysis methods must also be able to account for the multiple types of players involved. At minimum, they must model the traffickers, the victims, the buyers, and law enforcement. In some settings policy-makers, nongovernmental organizations, and other organizations would also be critical, as may community members themselves. Conceptually, game theoretic models can be extended to these settings, and some strides in this direction have been made in the terrorism risk analysis realm (Samuel & Guikema, 2012). However, traditional game theoretic approaches may not be ideal for this context given the strong assumptions underlying these models. Game theory tends to focus on individual agents as decision-makers, rather than collective entities such as legislative bodies and law enforcement (Cox, 2009). How can the interactions among the different groups and individuals be included in risk models in a way that reflects their different information states, decision-making, and behavior, and the fact that some are groups acting collectively rather than the individuals traditionally modeled in game theory? Again, more research is needed to advance risk analysis methods to better handle multiple players, particularly in light of the continuous nature of the problem and the behaviors being modeled.

2.5 Summary and Conclusion

It is clear from our research with anti-sex trafficking experts and law enforcement that they see a clear need for more advanced approaches and new frameworks to help them tackle this problem. Techniques and foundational knowledge from the risk analysis domain will not solve the sex trafficking problem by themselves. They can, however, be helpful in providing a holistic understanding of the problem to support decision-making and policy analysis. Additional research is needed, particularly in developing methods that can better address the continuous nature of the problem and the challenge of multiple interacting individuals and organizations. While these pose challenges, there is tremendous potential for risk analysis as a field to contribute in a significant way in the fight against sex trafficking. We must both adapt existing tools to this problem and develop the new models and frameworks needed.

Chapter 3

COVID-19 Pandemic's Impact on Online Sex Advertising and Sex Trafficking

Disruptive social events such as the COVID-19 pandemic can have significant impact on the commercial sex industry, yet these effects have been little understood. This paper examines the effect of the pandemic on one part of the commercial sex industry: sex trafficking through online advertising. Our analysis of a dataset of over one million sexual service advertisements scraped from the popular website rubratings.com suggests that the pandemic has had a significant effect on online advertising for sex in the United States, including advertisements we identified as likely being associated with trafficked sex workers. Ad volume decreased significantly around the time stay-at-home orders were put in place, then increased to levels well above prepandemic levels as COVID-related restrictions were relaxed. We contribute to the policy landscape by arguing that the initial decrease was associated with a loss of demand for sexual services due to pandemic-related health concerns, but that a confluence of factors, including the lack of economic and social support, increased the number of people vulnerable to being exploited. This study can assist policy-makers in predicting future changes in the sex industry to support a more just and inclusive society. In the context of future health crises, natural disasters, and major social disruptions, it can guide policy makers in apportioning public aid in a way that does not leave vulnerable populations and existing sex workers at greater risk of being trafficked.

Keywords: sex trafficking; online advertising; COVID-19 pandemic; commercial sex

Note: This research is currently under review with the *Journal of Human Trafficking*, submitted on February 12, 2021, and first reviews received on March 31, 2021.

Co-authors: Vanessa Castro, Bridgette Carr, Glen Bredin, and Seth Guikema

3.1 Background: Sex Trafficking and Social Disruption

Sex trafficking is a criminal enterprise in which a trafficker uses force, fraud, coercion, or age to exploit individuals to perform commercial sex acts. Sex trafficking is often depicted in mainstream media, and this media coverage plays an important role in the public understanding of sex trafficking. Examples include the 2019 U.S. media reporting on the involvement of well-known public figures such as Jeffrey Epstein in sex trafficking and the oft-repeated, but not evidence-supported, media claim that sex trafficking increases in response to singular events such as the Super Bowl. This type of coverage sensationalizes the crime and presents it as a response to specific, unusual events in a community. Such attention does not necessarily lead to positive outcomes (Ahmad & Sheshu, 2012; Caulkins et al., 2019; Lederer & Wetzel, 2014). Sex trafficking is an everyday problem that is widespread throughout the United States. The COVID-19 pandemic has the potential to substantially affect sex trafficking, and policy-makers need rigorous analysis of sex trafficking and its prevalence.

Sex trafficking primarily provides financial benefit to the traffickers and consequently traffickers must generate demand for the services of their victims. The primary way in which they do this is through online advertising (Boecking et al., 2019). Online advertising thus provides insight into the state of the sex trafficking in the United States. We focused on the subset of online advertisements with sex trafficking indicators: specifically, on indicators that the sex worker is managed by a third party.

Too often, research and reporting on sex trafficking focuses on unique social events (e.g., the Superbowl, the North American International Auto Show; Martin & Hill, 2019). Focusing only on one-off events reinforces the idea that sex trafficking happens only a few times a year or that it is a problem associated with the transportation of victims into a community and supported by an influx of buyers from outside that community. The term “trafficking” can be confusing in this context; sex trafficking is a criminal enterprise that exploits victims 365 days of the year and does not have to include movement. A victim of sex trafficking could be trafficked within their own home or neighborhood (Shively et al., 2012).

Rather than focusing on transportation of the victim, a better way to frame the idea of sex trafficking is to consider the presence of third-party management as a proxy. An indication of a third-party manager in an advertisement for commercial sexual services suggests that the provider is not an independent sex worker. While it is possible that an individual who is voluntarily participating in sex work may have a third-party manager, our informal interviews with law enforcement experts and others familiar with the sex trafficking industry suggest that this situation is rare[1]. Third-party management is thus strongly indicative of the individual being trafficked, as indicated in previous literature (Ibanez & Gazan, 2016).

While sex trafficking occurs in all facets of the commercial sex industry, this paper focuses on sex trafficking through online ads for in-person sex work. This distinction is important to understand the later claims we make about the connection between sex work and the COVID-19 pandemic. We recognize there is no one-size-fits-all portrayal of sex trafficking victims, as they can be found throughout the highly stratified and segmented commercial sex industry market,

[1] We are not attributing any feedback to a particular interviewee in order to honor their need to remain anonymous.

stretching from in-person street prostitution to brothels, massage parlors, clubs, exclusive escort services, and online virtual sex work (Shively et al., 2012). However, although not a direct measure of sex trafficking or available sex services, the volume of advertisements for in-person sexual services is a useful measure, as it helps characterize demand and supply (Ahmed & Seshu, 2012).

3.2 Hypothesis

While often-repeated media claims about that the presence of sex trafficking may lead one to believe sex trafficking is episodic and only associated with large events (Martin & Hill, 2019), sex trafficking is a persistent reality and outside events have an effect on its magnitude and profitability. Better understanding the relationship between sex trafficking and outside events could help multiple stakeholders such as law and policy makers and law enforcement to predict future changes and adapt to the changing nature of the industry. It can also help policy-makers better understand the effects—and limitations—of existing policies during times of social disruption.

We hypothesized that the onset of stay-at-home orders were associated with a dampening effect on commercial sex and sex trafficking online advertisements due to client and provider reluctance to contract the virus and a reduction in regular travel away from home for work and recreation on the part of potential clients. However, we also hypothesized that this temporary decrease would be accompanied by a confluence of other factors that would lead to a rapid increase in sex trafficking after restrictions lifted. These factors are detailed in the Discussion section. Prior to the discussion on factors, however, we first organize the article to detail the materials and methods used for our analysis. We test the hypotheses by collecting commercial sex advertising data scraped from a popular website, focusing on a subset of advertisements with strong indicators of trafficking and then testing for statistically significant changes in advertising volume.

In a recent study about plunging crime rates following government mandated stay-at-home orders, the authors warn against overgeneralizing crime and underestimating the specificity in offender decision-making (Stickle & Felson, 2020). In this research, we unpack this specificity to understand the nuanced changes the pandemic may have had on independent commercial sex work and sex trafficking. For example, temporal understanding is key in our analysis. Although we do not have data to support previous years, we draw conclusions about hypothesized decreases in online commercial sex advertising. We recognize that this may impact our analysis and that there exist many alternative hypotheses that would benefit from more data to improve accuracy. As Stickel and Felson (2020) suggest, the featured questions are as follows, "Does the pandemic have a measurable impact on online advertising associated with commercial sex and sex trafficking?" The logical corollary is, "If so, why have rates fallen?" and finally "What can be learned from this experience to leverage sex trafficking reduction in the future?"

3.3 Materials and methods

Building off other antitrafficking research (Boecking et al., 2019; Latonero, 2011; Nagpal et al., 2015), we used the volume of sex advertisements as an indicator of the prevalence of sex work in a given region.

3.3.1 Data Description

The internet hosts numerous online platforms for sexual services advertisements, which serve as a basis for providers to connect with buyers of commercial sex work. The landscape of available platforms has changed rapidly in the past year. Prior to the United States' legislative enactment of Fight Online Sex Trafficking Act—Stop Enabling Sex Trafficking Act (FOSTA-SESTA; Allow States and Victims to Fight Online Sex Trafficking Act, 2018), Backpage.com and Craigstlist.com had become the favored sites to facilitate sex work and sex trafficking (DeLateaur, 2016). After FOSTA-SESTA took effect, legislation effectively rendered Backpage.com's and Craigslist.com's personals advertisements ineffective on the internet for commercial sex advertising. Our informal expert interviews with law enforcement consistently inform us that a single replacement has not emerged to replace Backpage.com. Instead, websites spring up periodically and "hobbyists" and "mongers" (the terms used to describe frequenters of these websites) find their way to them by word of mouth (Feeney, 2013). However, in these conversations, one site that was consistently mentioned as a leading advertisement site was Rubratings.com.

Data used in this research consisted of over 1 million sex ads from the online platform Rubratings.com. This research is the first that we have been able to find that addresses sex trafficking anomaly detection in a post–FOSTA-SESTA era. We recognize that without one lead website that hosts the majority of advertisements to replace Backpage.com, we are analyzing only a portion of all existing online commercial sex advertisements online, potentially impacting our assessment. However, our interviewees indicated that despite the changes brought about by the FOSTA-SESTA era, online advertisements are still helpful indicators of the market for commercial sex in a given region.

The data from Rubratings.com were scraped in the Python programming language with the use of the BeautifulSoup package. This facilitated a regular automated scrape of the website from 140 cities across 47 U.S. states. The automated scrapers for this analysis retrieved 1,019,709 total advertisements in the United States over the period from January 3, 2020 to September 29, 2020. Data collected include paid advertisements used by clients to browse photos, services, rates, and other descriptions of sex workers. Typically, the advertisement will list an email address, a phone number, or both, in order for the client to make contact and set up a “date.” The collected data were the pageid, date modified, description text, location, phone number, and date scraped.

We were unable to obtain completely continuous data (i.e., every day) over the time period selected, as the Rubratings.com website and other similar websites are evolving and becoming savvier about how to block scraping. Over the 270 days between January 3, 2020 and September 29, 2020, there were 57 days on which the scraper did not run or was actively blocked. However, it is common for a given advertisement to be listed on the site for multiple days. This duplication means that most advertisements that ran on days the scraper was not running were caught by the scraper on a different day. The data for the remaining 213 days are thus likely representative of the entire time period under study. With this in mind, we removed the duplicated advertisements to accurately characterize the actual daily advertising behavior and to prevent inadvertently overweighting advertisements. We accomplished this by comparing the dates the advertiser posted the ad with the date the scraper captured the ad. If those dates matched we identified this as the advertisement’s unique appearance, appended that advertisement to the analytic sample, and deleted any identical copies. This procedure reduced the data set significantly.

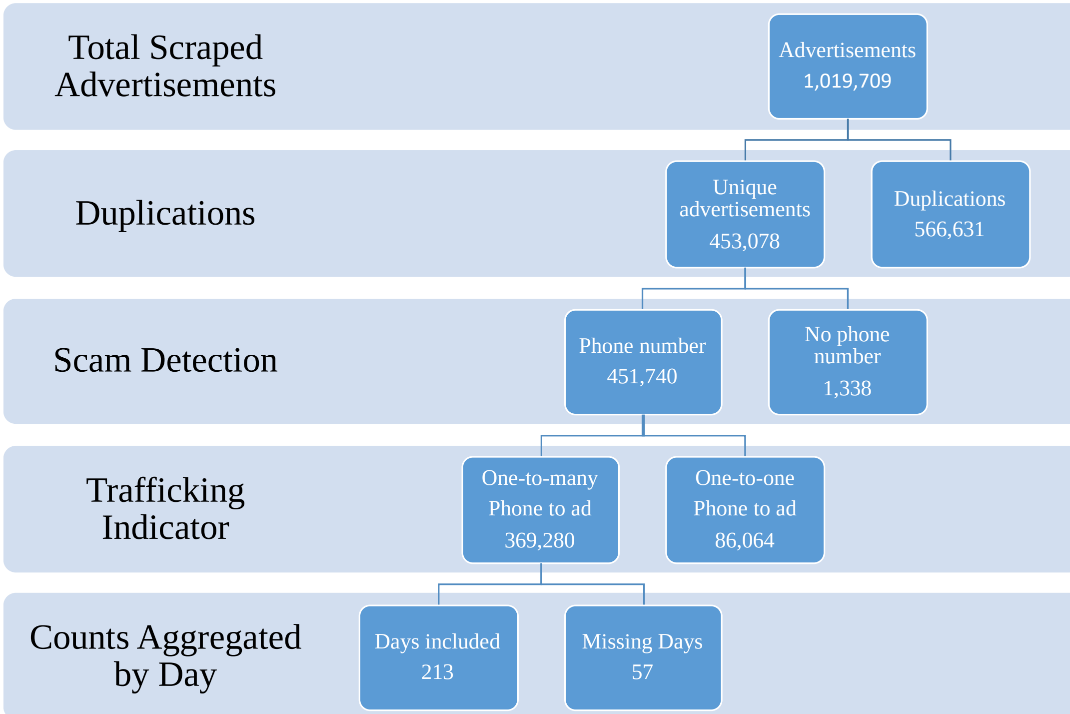

Figure 3.1: Data stratification for online advertising scraped from Rubratings.com from January 3, 2019 to September 29, 2020.

3.3.2 Data Stratification

Exploring the website and probing some advertisements quickly made it clear that many of these advertisements were actually scams attempting to obtain credit card information. Classified advertisements are fertile ground for scammers (Al-Rousan et al., 2020; McCormick & Eberle, 2013), and our informal interviews with anti-trafficking experts agreed that this is a common practice on these types of websites. One co-author (Guikema) confirmed this by sending emails to all advertisers in Ann Arbor, Michigan, who listed only an email address and found that all of them were scams attempting to steal credit card information. In related research in the context of spam detection, researchers found that 78% of fake or spam advertisements did not include phone numbers, while 87% of real advertisements did include phone numbers (Tran et al., 2011). For these reasons, the analytic sample includes only the advertisements with phone numbers for further analysis.

The data are also stratified to characterize the impact of sex trafficking separate from the impact of independent commercial sex work. Previous researchers based their characterization on a “new to town” stratification that labels advertisements that had recently appeared in that region (Boecking et al., 2019). We contend that this labeling is potentially misleading because it bounds sex trafficking to movement. As mentioned earlier, sex trafficking does not require movement, and this characterization may overlook the persistent risk to victims every day. Furthermore, this characterization ignores the reality that independent commercial sex workers move around and work in different cities as well. Instead of a “new to town” stratification, this research uses phone number clustering. Some advertisements point to the same phone number, whereas other phone numbers have single appearances in our data set. The informal interviews we conducted with law enforcement, current and former sex workers, and antitrafficking experts suggested that multiple advertisements pointing to one phone number are evidence of a call center or third-party management. These experts confirmed that the use of third-party management (or “pimps”) is a likely indicator of sex trafficking.

The idea of clustering advertisements by common parts to identify sex trafficking is not unique. Recent research on detecting sex trafficking in online advertisements hypothesizes that if the language in the advertisements have a lot in common, then there is reason to be suspicious (Aupperlee, 2021). Our research takes this concept a step further by using a common phone number as a reason for suspicious activity. This concept of phone number clustering is further reinforced by the informal interviews we conducted with current and former independent (nontrafficked) commercial sex workers, who described their phone number as their “lifeline” and a way to stay in touch with their "regulars," thereby suggesting that their singular advertisement would be connected only to their number. We incorporated this concept into our data by parsing the advertisements that had a relationship with only one phone number that did not appear to be related to any other advertisement (1:1 relationship between phone number and ad) from the clusters of many ads that pointed to one phone number (1:n, where n>1, relationship between phone number and ads). We recognize that some sex trafficking victims are likely using one phone number that will not appear elsewhere. We also recognize that the use of short-term phone numbers associated with a disposable phone or a "soft"/internet phone number is common in nefarious practices; this labeling will not capture those third-party managers who cycle through multiple disposable numbers (Chen et al., 2019). We do not wish to exclude these victims, but rather use this labeling as a generalized sex trafficking indicator. Figure 3.1 summarizes the stratification of the data set.

3.3.3 Data Limitations

Despite the richness of the data scraped from the online advertisements, there are admittedly many limitations as well. First, as mentioned previously, the automated scrapers did not capture advertising activity every day as intended. Instead, there were days when the scraper did not run or was actively blocked by the website. Despite regular retooling and restructuring of the scraper, some partial or entire days were not scraped; these were distributed at random throughout the entire data set. This observation is consistent with other similar research (Boecking et al., 2019).

Other limitations include the imperfect way the data are stratified for detection of scams and sex trafficking labeling. Research suggests that trafficking detection can be also done through analyzing the advertisement description itself through various types of machine-learning classification and by identifying movement patterns (Ibanez & Suthers, 2015; Nagpal et al., 2015; Szekely et al., 2015). Lastly, we recognize that law enforcement and antitrafficking nonprofit organizations place decoy advertisements on these types of websites to disrupt the criminal activity of those seeking paid sexual services. These decoys have the ability to add noise to the data, as they could be counted as either advertisements for commercial sex or for sex trafficking.

Lastly, our data is limited by its own temporal content and by its source. As discussed earlier, we recognize that our analysis would benefit from many more years of data in order to correctly address any yearly seasonality that may account for nuanced changes in advertising count. We also acknowledge that the shutdown of websites like Cityxguide.com, a site that gained popularity after the passage of the FOSTA-SESTA anti-sex trafficking legislation, could affect our analysis. Furthermore, some early media reporting suggest that many commercial sex workers chose to move online to websites like Onlyfans.com in the wake of stay-at-home orders (Boseley, 2020) which may also impact our analysis. Despite these acknowledged limitations, we believe this analysis can still be meaningful and contribute to a larger discussion of the pandemic's impact.

3.3.4 Methods

Following the data stratification, we were left with 369,280 advertisements that reflected indications of our sex trafficking indicators; see Table 3.1.

Table 3.1: Daily descriptive statistics for scraped advertisements (Rubratings.com) January 3, 2020 through September 29, 2020. Compiled, entire data set; Sex Trfk, records that exhibited the sex trafficking indicator; Comm Sex, records that did not exhibit the sex trafficking indicator, proxy for independent commercial sex workers

Data	U.S. Ad Count (total)	Minimum (by day)	Maximum (by day)	Mean (by day)	Std. Dev. (by day)
Compiled	1,019,709	210	8821	3711	1704
Sex Trfk	369,280	183	3194	1718	773
Comm Sex	86,064	93	883	406	274

We then aggregated these remaining advertisements by day and by each individual city. We normalized the number of daily commercial sex online advertisements by assessing an average number of advertisements in equilibrium (pre–COVID-19) and then dividing by that average over the hypothesized anomalous period. We also applied a 3-day moving average to smooth out the noise and emphasize the trends. We took the raw mean number of advertisements pre–COVID-19. Those entries were then divided by the pre–COVID-19 average to produce respective normalized values.

While data were collected from throughout the United States, we highlight eight major U.S. cities from a variety of geographic areas in order to demonstrate the varied impacts across the wide spectrum of stay-at-home order implementation dates. We chose to highlight the following regions based on the variety of stay-at-home implementation dates: Atlanta (April 3, 2019), Dallas/Fort Worth (April 2, 2019), Seattle (March 23, 2019), New York (March 22, 2019), Houston (April 2, 2019), Miami/Fort Lauderdale (April 3, 2020), and Detroit (April 3, 2019), as well as the entire United States (varied stay-at-home implementation dates; Kaiser Family Foundation, 2020). The process described above was repeated for all eight regions to obtain normalized values for each.

We systematically tested our hypothesis that there existed a decreasing trend of sex trafficking advertisements after stay-at-home orders were instituted by systematically comparing baseline and testing intervals around each individual state's stay-at-home orders (Kaiser Family Foundation, 2020). We tested the days immediately prior to the inflection point in 15-, 30-, 45-, and 60-day increments to the 15-, 45-, and 60-day increments immediately following the lockdown orders. We tested these pairings within the cities mentioned in the main test with a two-sided Welch's t test of the two independent samples. This statistical method tested the null hypothesis that the two samples have the same expected values. If the p value is smaller than .10, we reject the null hypothesis of equal expected values.

This testing helps us draw conclusions about the testing ranges' statistical significance. We are presenting a family of multiple hypothesis tests at the city level. We therefore make a Bonferroni correction at the city level and use a lower p value that accounts for the probability of getting one significant result from the family at the 10% level (Miller, 2012). We use the Bonferroni method because it is one of the most conservative and common ways to account for the number of comparisons. After using these methods to analyze records we categorized as showing indications of sex trafficking, we repeated the same methods for the records that were not classified by our approach as indicative of sex trafficking. We refer to these records later as a proxy for independent commercial sex workers. We conducted this same analysis for all eight cities as well as the United States as a whole.

3.4 Results

This section helps answer our first featured question from Section 3.2, "Does the pandemic have a measurable impact on online advertising associated with commercial sex and sex trafficking?"

Our results support our hypothesis that online ads flagged with our sex trafficking indicator would rapidly decline after stay-at-home orders began, and then rapidly increase after those orders were lifted. Figure 3.2 presents a time-series graph of total online commercial sex advertising volume and likely sex trafficking advertising volume for all regions studied.

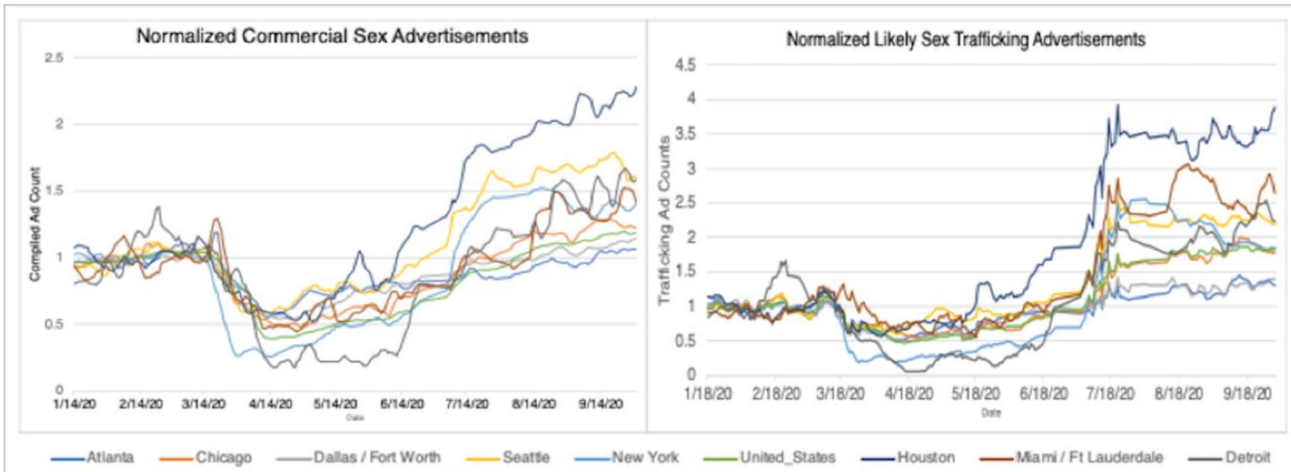

Figure 3.2: Number of advertisements for both compiled and "sex trafficking" data sets
*Note: Left (**2a**): Normalized time-series of daily ad counts of all compiled advertisements.*
*Right (**2b**): Normalized time-series of daily ad counts reflecting the sex trafficking indicator.*

It exhibits the number of daily online advertisements for both the raw data (Figure 3.2a) and the subset flagged as likely sex trafficking (Figure 3.2b), normalized by the pre-COVID advertising volume at the city level. Included in the graph are Atlanta, Chicago, Dallas/Fort Worth, Seattle, New York, Houston, Miami/Fort Lauderdale, and Detroit, as well as the United States as a whole.

Additionally, we list computed features on the trajectory and behavior of the likely sex trafficking advertisements for each city (Table 3.2). The table is sorted by the earliest stay-at-home order, with column descriptions listed (see Appendix A for methods of calculation). The table annotates the date at which that city's decrease in likely sex trafficking advertisements first began, followed by the average number of advertisements the city evidenced from January 1, 2020 until each city's decline began. We also calculated the velocity (approximate slope) at which each city's advertisements declined. The lowest number of advertisements at the bottom of each time-series, along with their corresponding dates, follow in the table; all of these figures assist in the calculation of the depth percentage of the decline.

Table 3.2: Statistics by city describing sex trafficking escort advertisement behavior over time. Period covering January 2020 to May 2020, sorted in order of stay-at-home orders.

city	stay_at_home	dip_began	pre_Covid_avg	dip_velocity	lowest_ad_count	lowest_ad_date	increase_began	depth (%)
Chicago	3/21/20	3/18/20	332.5	-6.24	164	4/14/20	4/30/20	49.32
New York	3/22/20	3/12/20	419.52	-19.53	107	3/28/20	4/14/20	25.51
Seattle	3/23/20	3/12/20	110.98	-2.37	66	3/31/20	4/20/20	59.47
Detroit	3/24/20	3/20/20	22.73	-0.72	4	4/15/20	6/11/20	17.60
Dallas / Fort Worth	4/2/20	3/17/20	606.69	-11.32	335	4/10/20	4/24/20	55.22
Houston	4/2/20	3/11/20	93.25	-1.71	54	4/3/20	5/5/20	57.91
Atlanta	4/3/20	3/16/20	333.21	-6.17	179	4/10/20	4/10/20	53.72
Miami / Ft Lauderdale	4/3/20	3/20/20	64.02	-1.67	29	4/10/20	4/27/20	45.30

We observed these statistics to be consistent with the visualized time-series graph in Figure 3.2. It is also worth noting that the dips in likely sex trafficking advertisements all began prior to the state-mandated stay-at-home order. This may have occurred in anticipation of government-required social distancing. Third-party managers of sex trafficked victims may have known that in-person sex work demand would decline in response, and were perhaps conserving the fees associated with posting advertisements online. In the early days of the outbreak in the United States, the risk communication around New York as the epicenter was clear (McKinley, 2020). We would expect that the response to the danger of the virus to be evidenced in a dramatic and rapid decline of in-person commercial sex, and that is indeed what we find in the data. Table 3.3 supports this characterization, as New York reflects the fastest velocity toward its lowest ad count.

Table 3.3 summarizes all of the statistical testing that characterizes the decrease in advertising during lockdown as well as the increase in advertising immediately following the lifting of restrictions for both the population of advertisements that we flagged as likely sex trafficking advertisements as well as those that we label as independent commercial sex workers.

Table 3.3: Summary of testing

City	All Ads (Stay at home)	All Ads (Lifted)	Flagged Ads (Stay at home)	Flagged Ads (Lifted)	Ind Sex Workers (Stay at home)	Ind Sex Workers (Lifted)
Atlanta	↓ (6)	— (4)	↓ (10)	↑ (8)	↓ (12)	↑ (12)
Chicago	↓ (9)	— (2)	↓ (11)	↑ (12)	↓ (12)	— (3)
Dallas/Fort Worth	— (3)	— (1)	↓ (9)	↑ (11)	↓ (12)	↑ (9)
Seattle	— (3)	↑ (7)	↓ (7)	↑ (12)	↓ (11)	— (4)
New York	↓ (12)	↑ (6)	↓ (12)	↑ (12)	↓ (12)	↑ (8)
Houston	— (2)	↑ (8)	— (3)	↑ (12)	↓ (9)	↑ (6)
Miami/Ft Lauderdale	↓ (12)	— (2)	↓ (6)	↑ (12)	↓ (9)	↑ (8)
Detroit	↓ (8)	— (2)	↓ (9)	↑ (12)	↓ (8)	— (4)
USA	↓ (6)	— (2)	↓ (12)	↑ (12)	↓ (11)	↑ (12)

Each cell represents twelve systematic pairings and the presence of an arrow represents statistical significance in six or more of those testing intervals, the number in parentheses indicates the number of pairings that were statistically significant, and the direction of the arrow indicates an increasing or decreasing trend.

3.4.1 Decrease in Advertising During Lockdown

A statistically significant decrease is evident in both compiled advertisements and likely sex trafficking advertisements immediately after stay-at-home orders were issued in most but not all cities. The trend of decreased sex advertising during lockdown was consistent for the United States as a whole and specific cities. Taking New York as an example (stay-at-home order on March 22, 2020), we found a pre-COVID mean number of advertisements of 419.52 (see Table 3.2). The first four entries for the 3-day moving average for New York were 396, 374.4, 350.2, and 331.4. Those entries were then divided by the pre–COVID-19 average to produce their respective normalized values of 0.944, 0.892, 0.835, and 0.790.

In the New York example, a *p* value greater than .0087 would indicate that there exists no significant difference in the expected values of the baseline and testing ranges. In Table 3.4, we see that New York's data show that all range comparisons reveal statistical anomalies. Furthermore, the negative *t* statistic indicates that the advertising rate decreased. We can confidently conclude that New York experienced a downward trend of likely sex trafficking advertisements associated with stay-at-home orders. One can also observe that the decrease began prior to the state-mandated stay-at-home order. We could surmise that advertisements started to drop in preparation for, or perhaps in anticipation of, the orders.

Table 3.4: New York City intervals before and after stay-at-home order (March 22, 2020).

city	pre-covid	baseline interval (days)	stay_at_home	testing interval (days)	t_stat	p_val
New York	['2020-03-06', '2020-03-21']	15	['2020-03-22', '2020-04-06']	15	-4.770067388	0.0005219
New York	['2020-03-06', '2020-03-21']	15	['2020-03-22', '2020-05-06']	45	-4.675141507	0.000640865
New York	['2020-03-06', '2020-03-21']	15	['2020-03-22', '2020-05-21']	60	-4.374609426	0.001038278
New York	['2020-02-20', '2020-03-21']	30	['2020-03-22', '2020-04-06']	15	-8.211747406	6.86E-08
New York	['2020-02-20', '2020-03-21']	30	['2020-03-22', '2020-05-06']	45	-8.122528606	1.02E-07
New York	['2020-02-20', '2020-03-21']	30	['2020-03-22', '2020-05-21']	60	-7.704002297	2.09E-07
New York	['2019-02-05', '2020-03-21']	45	['2020-03-22', '2020-04-06']	15	-23.85485296	1.67E-46
New York	['2019-02-05', '2020-03-21']	45	['2020-03-22', '2020-05-06']	45	-25.19075404	3.85E-65
New York	['2019-02-05', '2020-03-21']	45	['2020-03-22', '2020-05-21']	60	-23.2173616	2.24E-61
New York	['2019-01-21', '2020-03-21']	60	['2020-03-22', '2020-04-06']	15	-23.85485296	1.67E-46
New York	['2019-01-21', '2020-03-21']	60	['2020-03-22', '2020-05-06']	45	-25.19075404	3.85E-65
New York	['2019-01-21', '2020-03-21']	60	['2020-03-22', '2020-05-21']	60	-23.2173616	2.24E-61

We found statistically anomalous decreases in likely sex trafficking advertisement in Atlanta, Chicago, Dallas/Fort Worth, Seattle, Miami, Detroit, New York, and the aggregated data across the United States. As Table 3.5 exhibits, however, most of the tested ranges for Houston evidenced *p* values greater than our critical value. This finding implies that Houston did not experience a statistically significant dip in likely sex trafficking advertisement after their stay-at-home order. The result is consistent with media reporting about generalized non-compliance with stay-at-home orders and the accelerated or constant growth rates of new coronavirus cases in those regions (Advisory Board, 2020).

Table 3.5: Houston, Texas intervals before and after stay-at-home order (April 4, 2020)

city	pre-covid	baseline interval (days)	stay_at_home	testing interval (days)	t_stat	p_val
Houston	['2020-03-17', '2020-04-01']	15	['2020-04-02', '2020-04-17']	15	0.169146877	0.866934077
Houston	['2020-03-17', '2020-04-01']	15	['2020-04-02', '2020-05-17']	45	1.692157446	0.104854247
Houston	['2020-03-17', '2020-04-01']	15	['2020-04-02', '2020-06-01']	60	2.96437709	0.006572771
Houston	['2020-03-02', '2020-04-01']	30	['2020-04-02', '2020-04-17']	15	-1.561868154	0.126138093
Houston	['2020-03-02', '2020-04-01']	30	['2020-04-02', '2020-05-17']	45	-0.317110334	0.753009709
Houston	['2020-03-02', '2020-04-01']	30	['2020-04-02', '2020-06-01']	60	0.941790546	0.351941951
Houston	['2019-02-16', '2020-04-01']	45	['2020-04-02', '2020-04-17']	15	-0.435539923	0.6667978
Houston	['2019-02-16', '2020-04-01']	45	['2020-04-02', '2020-05-17']	45	2.116431888	0.03640258
Houston	['2019-02-16', '2020-04-01']	45	['2020-04-02', '2020-06-01']	60	4.123370066	6.73E-05
Houston	['2019-02-01', '2020-04-01']	60	['2020-04-02', '2020-04-17']	15	-0.435539923	0.6667978
Houston	['2019-02-01', '2020-04-01']	60	['2020-04-02', '2020-05-17']	45	2.116431888	0.03640258
Houston	['2019-02-01', '2020-04-01']	60	['2020-04-02', '2020-06-01']	60	4.123370066	6.73E-05

3.4.2 Increase in Advertising After Lockdown

After lockdown restrictions began to ease, online advertising quickly rose again. Likely sex trafficking advertisements increased consistently in all regions studied. However, we found differing results for likely independent commercial sex workers across the different cities. Table 3.3 and Figure 3.2 show that in the period following the lifting of restrictions, the change in volume for advertisements likely associated with independent sex workers are mixed, with some rising rapidly in July (New York, Detroit) while others rise gradually (Dallas) or decline soon after the lifting of restrictions (Miami). By contrast, likely sex trafficking advertising is seen to rise rapidly in July across all of our tested cities. This discrepancy may be an indication that traffickers are forcing their victims into dangerous client interactions as soon as possible to drive income, in contrast to their independent peers who may be able to choose to be more risk-averse, waiting to re-enter the market until pandemic conditions improve.

Table 3.3 shows the results of statistical comparisons of prepandemic advertising to two postpandemic time periods: immediately after stay-at-home orders were issued, and in the longer term after stay-at-home orders were lifted. In addition to the statistically significant decrease associated with the first stay-at-home orders discussed earlier, we also see a statistically significant increase in advertisements with trafficking indicators in all cities when comparing the pre-COVID period with the period immediately following lifted restrictions. From Figure 3.2b we see that this advertising increased from between 1.25 and 3.5 times pre-COVID levels, depending on the city.

The substantial and statistically significant increase in sex trafficking advertising could be explained by an increase in the number of sex trafficking victims post-COVID. Alternately, traffickers experiencing significant declines in income due to the pandemic could be attempting to make up for this through increased advertising. As will be discussed below, we believe the former reason may account for some of the rise in advertising, as the increased economic precarity associated with the pandemic may make more people vulnerable to being trafficked.

3.5 Discussion

3.5.1 Reasons for Trends in Sex Advertising

This section addresses our second featured question from Section 3.2, “If our hypothesis is true, why have rates fallen?” The present research confirms the hypothesis. In addition, the observed discrepancy between data for ads we identified as leading to independent sex workers versus likely sex trafficking victims show a clear difference, indicating that the chosen method of identifying ads linked to trafficking victims was effective.

As discussed above, we would expect the COVID-19 pandemic to lead to a temporary decrease in both independent commercial sex work and sex trafficking. Under normal conditions, an overall decrease could indicate the favorable outcome of fewer sex trafficking victims. The pandemic’s confluence with a variety of socioeconomic, environmental, and policy factors, however, creates a potential wave of sex trafficking victims that should not be disregarded by government policy and decision-makers. COVID-19 will have outsized impacts on a marginalized population whose vulnerability is amplified by the pandemic, making them more likely to enter a trafficking situation, as well as a population of existing sex trafficking victims who are more likely to remain in exploited conditions as a result of the pandemic. A proactive pandemic response needs concrete policy actions to prevent a likely increase in sex trafficking exploitation. Because sex trafficking by its very nature affects multiple sociopolitical and economic systems and the health of at-risk communities, further analysis of the push and pull factors provides a framework for the estimation of how COVID-19 may affect the commercial sex industry and sex trafficking (Worsnop, 2019).

3.5.2 Socioeconomic Factors

The onset of the pandemic and subsequent social-distancing orders makes in-person commercial sex work a risky occupation, placing sex workers and sex trafficking victims in an even more economically insecure state. Furthermore, the economic aid provided by the U.S. government through the Coronavirus Aid, Relief, and Economic Security (CARES) Act is subject to anti-sex work measures built into the eligibility provisions, as the federal regulations determining which businesses may receive such aid overtly discriminate against businesses that derive revenue from services' depictions or displays of a "prurient sexual nature" (Code of Federal Regulations, 2017). This should, over time, make some otherwise independent commercial sex workers more vulnerable to becoming trafficking victims due to increased socioeconomic vulnerability. At the same time, those who have lost jobs due to COVID-19 and work in the gig economy may not have access to unemployment insurance as they are not a part of the formal economy. Qualifying for such benefits will therefore likely be challenging for these individuals and, subsequently, make these individuals more susceptible to sex trafficking. These factors would tend to suggest that the increase in sex trafficking advertising is at least partially due to an increase in the number of sex trafficking victims available to traffickers, allowing them to increase their operations postpandemic.

3.5.3 Environmental Factors

The effects of COVID-19 are most profoundly felt in marginalized communities experiencing heightened vulnerability due to factors such as race, income, gender, and housing instability (Smith & Cockayne, 2020). Consequently, during these unprecedented times, these factors may increase the exploitation of people in vulnerable positions. The U.S. Department of State reports that although sex trafficking affects every demographic, a common factor is the victim's vulnerability to exploitation. Traffickers are able to thrive when inequality exists and where people lack access to social protection and justice (Department of State, 2016). The pandemic creates unique environmental conditions that further restrict the access these already marginalized communities have to social services and amplifies their existing risk of sex trafficking (Todres & Diaz, 2020).

The pandemic intensifies this marginalization, making individuals more likely to stay in a trafficked situation or become vulnerable to sex trafficking. For example, efforts to establish preventive measures to quell the spread of the coronavirus have led to the displacement of already marginalized commercial sex workers in brothels, thereby causing them to become homeless during this crisis (Berger, 2020). Similarly, as federal stimulus unemployment benefits and eviction bans expire across the United States, evictions will likely increase, making vulnerable individuals homeless (Cowin et al., 2020). Further compounding their inequality and lack of a safety net in both cases, this displacement cannot be addressed by conventional solutions because of the unique pandemic environmental conditions that restrict access to essential social services. This confluence of conditions has the potential to keep existing victims in dangerous sex trafficking situations or place others at risk of entering sex trafficking. If governments do not address the potential effects of the pandemic on marginalized populations, we are likely to see a sustained rise in sex trafficking prevalence.

3.5.4 Government Efforts and Policy Factors

In addition to policies that continue to stigmatize sex workers during COVID-19, it is important to consider the ways in which public health measures to contain the virus through contact tracing may further exclude sex workers and sex trafficking victims from protection. In most of the United States, because commercial sex work is criminalized, there may be gaps and underreporting in surveillance and tracing, as compliance may jeopardize sex workers' ability to maintain their privacy and client confidentiality. Policing may further threaten the health and financial well-being of sex workers, who may face lost wages or clientele not willing to participate in public health surveillance efforts that may require documentation and reporting. When stay-at-home restrictions lift, it stands to reason that sex workers and traffickers will redouble their efforts to make up for lost wages. In this way, the criminalization of sex work may have an inadvertent influence on increased postpandemic commercial sex and sex trafficking levels.

One way to contain the spread of the virus is to implement policy measures that will not criminalize sex workers, while continuing to advocate for criminalization of all other parts of the sex trafficking enterprise to protect victims. Decriminalization of sex work would not only eliminate the burden of police interactions, but also allow commercial sex workers and sex trafficking victims to qualify for economic relief currently reserved for members of the formal economy. A full decriminalization strategy may face insurmountable legislative constraints, but even temporarily reinforced social services and local policing strategies that limit interactions with sex workers would help keep them from becoming further marginalized and vulnerable to victimization.

Finally, we address our last featured question, "What can be learned from this experience to leverage sex trafficking reduction in the future?" Our data analysis suggests that the rapid increase in sex trafficking that we hypothesized would occur based on socioeconomic, environmental, policy, and regulatory factors is already happening. Sex trafficking advertising has increased substantially above prepandemic levels, and these increases are statistically significant. In order to learn from this experience to apply them to any future pandemic or major social disruption, people in positions to influence antitrafficking efforts should recognize the outsized impact these events are having on vulnerable populations and sex trafficking victims. Some possible preparations include reinforcing social services, making concrete steps of inclusion for marginalized populations, and even considering a path towards decriminalization of sex work to give independent commercial sex workers and sex trafficking victims a safety net. We urgently need action to prevent this increase from accelerating even more rapidly.

Chapter 4

Time-Series Anomaly Detection Algorithm Comparison for Risk-Based Decision-Making

Time-series anomaly detection has many applications across disciplines such as healthcare, manufacturing, quality control, finance, cybersecurity, and crime prevention. In this article we characterize the effectiveness of both traditional statistical approaches in addition to a machine-learning approach and evaluate their effectiveness in risk-based decision-making. We provide this collection of techniques to advance the univariate time-series anomaly detection literature by drawing from concrete examples of biosurveillance, manufacturing, and quality control to generalize their detection power and robustness for the risk analysis community. We compare their associated machine learning, statistical process control, and change point detection techniques and provide a reliable comparison using a synthetic data set that contains point, collective, and contextual anomalies (i.e., those that involve a single unusual data point, series of unusual data points, and data point(s) that are identifiable as unusual only with knowledge of some aspect of their context). We frame the analysis using a broad spectrum of approaches to examine the hypothesis that algorithms drawn from multiple disciplines will perform differently for varying types of anomalies. Our results show that the machine learning approach performed the most consistently across all three different types of algorithms; whereas the change point detection was best at discovering the contextual anomaly and the statistical process control approach was superior to others at detecting the collective anomalies. Finally, we demonstrate how these tools can be used as a risk-based approach to fighting sex trafficking by detecting anomalies in online advertisements.

Keywords: time-series, anomaly detection, machine learning, change point detection, statistical process control, sex trafficking, online advertising

Note: This research will be submitted to *Risk Analysis*.

Co-author: Seth Guikema

4.1 Introduction

The volume and ubiquity of data makes it difficult to find a signal in the noise. Anomaly detection, a technique of finding unusual patterns or activity relative to a standard point, can be an important tool in discovering rare events among a series of observations over time. This concept has a wide variety of applications in business, manufacturing, cybersecurity, fraud detection, healthcare, sensor monitoring, and many more diverse domains (Lane & Brodley, 1997; Laptev et al., 2015; Parpoula et al., 2017). The research community has become keenly interested in anomaly detection due to its wide range of applications and the vast amounts of data streaming in from sensors (Fahim & Sillitti, 2019). An anomaly is defined by most scholars as an abnormal data point (or series of data points): one that differs significantly from the surrounding data, in a way that cannot be attributed to random error (Wu, 2016). The normality or abnormality of a piece of data depends on a variety of contextual factors; Fahim and Sillitti (2019) give the example of a spike in traffic being considered normal during working hours but anomalous at midnight.

Data that can be used to detect anomalies are collected by machines, sensors, regulators, and websites and may be collected continually on a per-hour, per-minute, or per-second basis for long periods of time. The temporal component of this enormous amount of data adds a layer of complexity to analysis. In business decision-making, the concept of time is acknowledged as an integral component of risk-based analysis. For example, stock prices are analyzed, modeled, and forecasted as a function of time, with analysts seeking to identify chronological trends and predict events in the near future (Wu, 2016). Time-series risk management models were central to some debates between the finance industry and regulators after the 2008 financial crisis, with industry professionals arguing that accepted models failed to predict extreme and volatile stock market risk exposure (Kim et al., 2011). Along the same lines, some research shows how tightly coupled climate change and insurance markets are by using a univariate time-series approach (Romilly, 2007). Risk researchers similarly deal with time-series data regularly and may benefit from this research's specific examination of the intersection of anomaly detection and time-series data as it pertains to risk analysis.

Such techniques have shown effectiveness in a wide variety of fields. For example, biosurveillance systems designed to detect epidemic outbreaks use sophisticated modelling techniques to notify decision-makers of public health threats (Fricker, 2013). These techniques use data on such factors as chief complaints of emergency room patients to identify a possible outbreak of a disease such as H1N1 before confirmed cases have begun to appear. In the realm of commercial aviation, researchers have used time-series data on factors including GDP, air travel passenger turnover, and total number of airline employees to identify influencing factors and provide a theoretical basis for understanding high-risk events, including accidents, in airport flight areas (Shao et al., 2020). Decrouez and Robinson (2013) used time-series risk analysis to dynamically classify material for border inspection, "profiling" certain shipments as likely or unlikely to carry pests based on how often shipments of the same product or from the same region have carried pests within a recent time period. Li et al. (2017) used data from the Global Terrorism Database to predict the conditional probability of bombing attacks in Iraq and Afghanistan while comparing these data to those on armed assaults (Li et al., 2017). Time-series anomaly detection in terms of security and risk management (credit risk, fraud detection; Ahmed et al., 2016) is already a well-researched area, but there exist many unrelated areas of study that can be fruitfully examined by risk analysts in the future. As we will discuss later, it is unclear how traditional risk analysis models or anomaly detection would apply to sex trafficking, or more generally to noisy online advertisements. This research extends the literature by comparing the robustness and effectiveness of various approaches when used on sex trafficking data under different conditions.

In a recent paper celebrating the anniversary of the Society for Risk Analysis, an esteemed group of risk researchers cataloged and highlighted the accomplishments of the risk analysis community over the last 40 years (Greenberg et al., 2020). They also presented a list of the most salient risk-related challenges for the next decade. Among those included were another future pandemic, climate change, the impact of automation in industries such as agriculture, cascading disasters in infrastructure, and anthropogenic hazards such as terrorism and war. All these topics have an inherent temporal component, and analyzing these challenges will likely involve contextualizing and analyzing the relevant data with a time-series approach. For example, another major pandemic will produce streams of disease prevalence data and statistics, as we have seen throughout the COVID-19 pandemic. As in years past, a future war of large magnitude would place many analysts and decision-makers in a position of trying to make sense of many kinds of time-series data, whether it be the number of bombing attacks, armed assaults, improvised explosive devices, or perhaps a series of satellite images used to make strategic decisions. Risk researchers are no strangers to extreme events, having written prolifically about war, natural disaster, and economic depressions (Greenberg et al., 2020), and that will not likely change in the next decade. The purpose of the research in this chapter is two-fold: 1) to characterize the effectiveness of anomaly detection algorithms generalized for risk-based decision-makers, and 2) to contribute a case study using sex trafficking prevalence data applied to the researched techniques. As we brace for the future, this paper will offer suggestions for potential use of time-series anomaly detection for risk-based decisions. It will discuss the type of anomalies to be used in such analysis, and also discuss three anomaly detection techniques for possible further implementation in the risk analysis community: machine learning, change point detection, and statistical process control. Using a synthetic data set, we will present methods for data preprocessing and algorithm comparison. All three approaches will be evaluated against key metrics that measure the effectiveness and detection power of each method. Lastly, we will close out the analysis with a use case on a real-world temporal data set that counts the number of suspected online sex trafficking advertisements in order to demonstrate how well each algorithm performed and recommend ways in which policy-makers can incorporate these methods into decision-making. These insights will show how temporal anomaly detection can provide an improved basis for risk-based decision-making.

4.2 Background

4.2.1 Types of Time-Series Anomalies

The ability to recognize anomalies can have enormous effects on a system. In a sensor-driven environment, an anomaly signaling alarm can be an indication of wasted resources or perhaps even a security intrusion (Fahim & Sillitti, 2019). In a comparison of unsupervised anomaly detection algorithms, Goldstein and Uchida (2016) divide anomalies into three types: point anomalies, collective anomalies, and contextual anomalies . A point anomaly occurs in a time-series when a single data point differs significantly from the rest of the series, better known as an outlier. For instance, if a series of daily temperature readings includes a single day that is 10 degrees hotter than any other day in a 60-day period, that temperature reading would be a point anomaly. Collective anomalies, on the other hand, are anomalies that are temporally grouped as a set of instances. For example, high electricity consumption over the course of a several-day heat wave might be considered a collective anomaly. Lastly, a contextual anomaly is a data point or series of points that could be perceived as normal until some kind of postprocessing context-aware technique or analyst with domain expertise detects an anomaly. For example, a high volume of highway traffic on the Wednesday before Thanksgiving might appear anomalous if the data were compared to those from other Wednesdays, but would appear normal to analysts taking into account that many families drive to visit families for the Thursday holiday (Hayes & Capretz, 2014). This type of detection involves deep knowledge of the data and their context. In this work, we compare three different anomaly detection algorithms for point, collective, and contextual anomalies.

4.2.2 Common Anomaly Detection Approaches

We present these three methods as a means for researchers to compare and contrast different forms of univariate time-series anomaly detection for risk-based decision-making. In the domains often subjected to risk analysis, data have three common important features: 1) they are temporal, 2) the estimations of risk are often drawn from one observation or variable (such as temperature, water level, or volume of traffic per hour), and 3) the events they are trying to prevent are relatively rare. For example, wildfire researchers employ risk-based decision-making to develop optimal strategies for community risk mitigation and resource allocation (Ager et al., 2015; Wilson et al., 2011). Similar to the data sets we use as examples, these events produce temporal count data that are often imbalanced because wildfires are not a frequent occurrence. Another area in which researchers apply risk management principles is occupational safety; they analyze data of workplace accidents to highlight environmental factors that contribute to occupational injuries (Galizzi & Tempesti, 2015). Further exploration into this type of analysis could include anomaly detection methods on their zero-inflated models. In the past, risk analysts have empowered other researchers to use their tools for reducing risk of complex problems such as terrorism (Deisler, 2002). As a further example of future applications of these methods, we can enrich existing work in adversarial risk analysis in the context of resource allocation strategies for law enforcement (Gil et al., 2016; Huddleston et al., 2008). By using time-series anomaly detection on crime statistics, we can better allocate law enforcement resources. These types of techniques applied to risk-related challenges could help analysts decide where to focus attention.

4.2.2.1 Machine Learning

In machine learning, the pursuit of abnormal conditions has always been of interest to researchers (Goldstein & Uchida, 2016). Machine learning methods are numerous, including approaches such as regressions, random forests, neural networks, and support vector machines, while the number of application domains is even greater (Breiman, 2001; Friedman et al., 2001). The key strength machine learning approaches bring to anomaly detection, however, is their ability to predict a value and compare it to the observed value. Large prediction errors can often signal an anomaly. The COVID-19 pandemic, threats of future pandemics, and bioterrorism present opportunities to leverage the strengths of machine learning to prevent catastrophic harm. Early identification of these threats through biosurveillance powered by machine learning can provoke epidemiological investigation and enable a public health response to mitigate the risks (Fricker, 2013). Although the practice of biosurveillance involves a variety of statistical methods to detect disease outbreak, we use the term “biosurveillance” as shorthand to describe a family of machine-learning approaches of time-series anomaly detection in the realm of infectious disease. We do this because biosurveillance methods have received a good deal of attention from the research community and because this approach aligns well with our goal of evaluating time-series anomaly detection algorithms for the risk analysis community. We recognize that these methods are not confined to use within biosurveillance and are also applied to demand estimation, sales forecasting, and many more applications. This approach will often use time-series forecasting and error measurements between predicted and actual values to determine if an anomaly exists (De Gooijer & Hyndman, 2006). Similar to many machine-learning approaches, forecasting in a biosurveillance context will generally consist of three steps: parameterization (collecting raw data from a monitored environment), training, and detection (Omar et al., 2013). Using temporal lags as variables, researchers may then apply a machine learning model to create predictions that they can then evaluate against the testing set of data.

4.2.2.2 Change Point Detection

Although change point detection has its roots in quality control, it sees uses today in settings ranging from intrusion detection to climate change. Change points detection identifies discrete, abrupt transition points in the time-series structural properties, such as mean and variance (Montgomery, 2012). In a quality-control setting, change point detection is especially effective for monitoring the quality of a process as sustained shifts occur in that process’ product. This approach determines whether the underlying model has undergone a distributional change of condition; for instance, detecting changes in physiological data to identify disease, or detecting changes in temperature to determine whether climate change has taken place (Aminikhanghahi & Cook, 2017).

One well-established method for change point detection is binary segmentation, in which an entire time-series is analyzed with a single change point method to detect whether a change in variance exists (Killick et al., 2012). In binary segmentation, a data analysis method identifies the most significant change point in a temporal sequence of data (for instance, the largest rise or fall in temperature in a series of daily temperature readings) and splits the series in two around this change point. The same procedure is then used to split the resulting subseries in two, and so on. This process continues until no further change point splits are detected in the subsets (Rohrbeck, 2013). This method is low in complexity and can allow researchers to detect multiple change points using a single change point detection method.

4.2.2.3 Statistical Process Control

Statistical process control aims to eliminate variability around production goals, and often employs statistical techniques to control the production method. The most common statistical process control method is the control chart (Shewart, 1931). Although a major objective of a control chart is to find assignable causes of process shifts in order to identify and reduce variability, the concept can be applied to most forms of univariate time-series data. The estimated mean and standard error provide a basis for establishing thresholds that can be used to compare testing data with a distribution that reflects a “normal” state. Points that exceed these thresholds are labeled “rare” and trigger a signal to indicate an anomaly. In assigning these thresholds, researchers must balance the goals of correctly detecting and labeling anomalies and of ensuring a relatively low rate of false signals so they do not require constant verification and investigation (Fricker, 2013). No manufacturing process remains under control forever or always follows the same distribution, and the same could be said of other time-series applications of statistical process control. For example, the risk-based use case we will study later examines the daily count of online commercial sex work advertisements; it evidences anomalies that would be considered "out of control" from a manufacturing perspective.

4.2.2.4 Online vs. Offline Anomaly Detection

All three of these methods can be used in an "online" or "offline" configuration. In the former, the researcher is either looking at the data forensically—examining monitored data in retrospect—or analyzing a given, closed body of data. In the former, the algorithm takes advantage of new or incoming data as they stream in. The present research will focus primarily on offline approaches, also known as retrospective or *a posteriori*. They take advantage of an entire data set to detect anomalies, and tend to be more accurate than online approaches that incorporate only the most recent information (Fricker, 2013). In the biosurveillance domain, for example, a machine learning model would typically be trained only once and then used continuously for prediction, a process that would not incorporate streaming data. We are also sensitive to the needs of risk analysts who may benefit from this work's algorithm comparison, and their recurring work with government agencies whose data are frequently provided in an offline format.

4.3 Materials and Methods

This section will detail the experimental methodology we used in order to systematically compare three different algorithms for univariate time-series anomaly detection. Our three algorithms were chosen as representative, concrete examples of the broader classes they come from. We start with a data set that is illustrative of a risk-based decision-making approach: anomaly detection for sex trafficking. As previously discussed, risk-related data are often temporal and univariate, and anomalies are rare, making the real-world sex trafficking data set conducive to analysis. We model our synthetic data sets to match general characteristics of the real-world data set in order to compare our three methods of anomaly detection. Synthetically generated data allows us to build a risk-informed analytic sample with artificially added known anomalies for detection and allows for more extensive testing. Figure 4.1 summarizes the overall methodology, starting with the preprocessing of our data that produce our analytic sample. This sample is what we use to generate our synthetic data that we treat with the machine learning, change point detection, and statistical process control algorithms. These approaches produce classifications that we use to evaluate the results of their performance. We consolidate our findings and apply them to a real-world use case of anomaly detection for sex trafficking. This section expands on each one of the below steps.

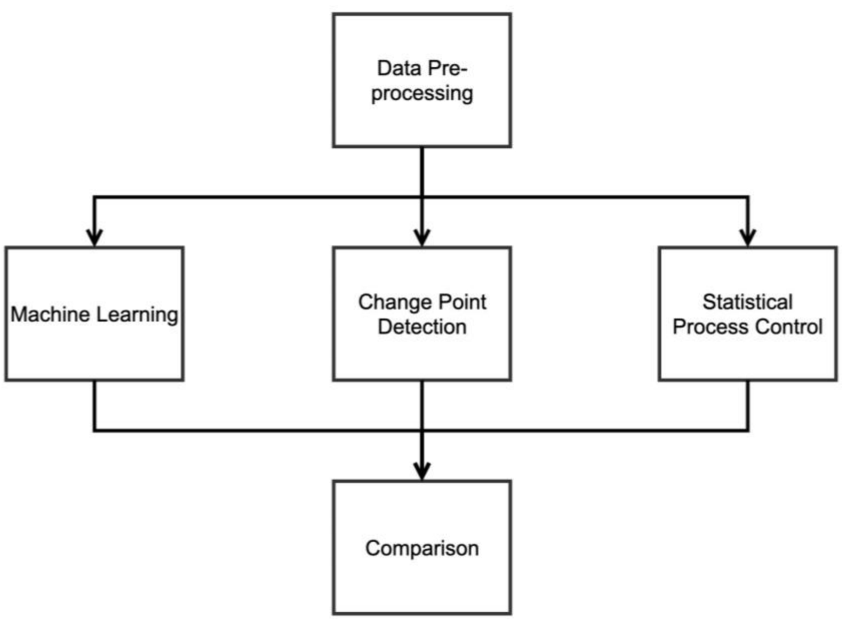

Figure 4.1: Overview of methodology

4.3.1 Materials

4.3.1.1 Real-World Data Description

The original data set will be covered in more detail in the case study and in Appendix B, but in order to contextualize the synthetic data set we provide the preprocessing and transformation steps used to produce the analytic sample. We model the synthetic data sets we use in the algorithm comparison on the analytic sample. In order to acquire the real-world data, we build on existing research that uses online advertising as a proxy for sex trafficking and scrape advertisements to provide count data across the United States from July 11, 2019 to May 31, 2020. This period of time allows us to investigate a hypothesized anomalous period after the initial stay-at-home orders that the United States implemented across the country in response to COVID-19.

4.3.1.2 Data Preprocessing

Data often need to be examined and modified by analysts before being processed by an algorithm. Nielsen (2016) summarizes best practices for time-series data preparation, and recommends the following steps:

1. Check for stationarity
2. Check for seasonal and trend influences
3. Correct for stationarity if necessary

These three steps are important for enabling the use of past data to make accurate predictions of future data. Generally, we want the data to continue to follow the same general trends over time. Stationary time-series data have the same characteristics across every time interval (Braei &

Wagner, 2020). A verification of stationarity is necessary because most time-series models assume that the underlying data have a constant mean/variance and that the autocovariance does not depend on time. This feature in turn allows us to leverage the statistical properties of certain models, such as machine learning, for forecasting (Nielsen, 2016). Research also suggests that using stationary data makes it possible to construct one model for anomaly detection instead of having to create multiple models to account for changes in the series (Talagala et al., 2020). This paper uses a series of quantitative and qualitative techniques to check for stationarity: 1) Split data, check mean and variance; 2) Autocorrelation check; and 3) Augmented Dickey-Fuller test (see Appendix B; Mushtaq, 2011).

All three tests indicated lack of stationarity, an outcome that makes it necessary to transform the data in order to make them stationary. The real-world time-series evidences two different means and variances for the two halves of the data. While this is not a clear-cut indication that stationarity is not present, the result was confirmed by the autocorrelation check and Augmented Dickey-Fuller test, which suggested that the time-series has a unit root and is not stationary (see Appendix B). It is common for time-series data to be nonstationary because of trends and seasonality; we can correct for this through the use of log transformations, smoothing, differencing, polynomial fitting, and decompositions (Nielsen, 2016). Failing to do so could result in mis-specified variance, poor model fit, or a failure to leverage the time-dependent nature of the data. We also experimented with the integration of trends and seasonality into the data and found that the addition of each component resulted in increased errors, which led us to conclude that their inclusion would not improve model performance (see Appendix B).

A common way to conduct time-series transformation is by using differencing and logarithmic (log) transformations (Fricker, 2013). Differencing is one way to make a nonstationary time-series stationary by computing the differences between consecutive observations. Additionally, we added a log transformation to help stabilize the variance of the time-series. A visualization of the differencing and the log transformation appears in Figure 4.2. This transformation lends itself well for analysis, as the data become more uniform throughout the time-series. This vector of counts over time serves as the basis for creating the synthetic data set.

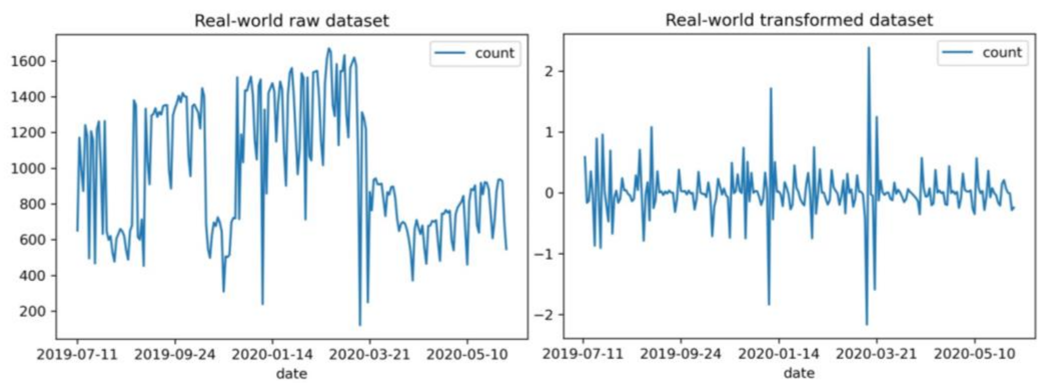

Figure 4.2: Raw data aggregated advertising counts and differenced-log transformed data

4.3.1.3 Synthetic data set generation

We model our synthetic data set to resemble the real-world data set to demonstrate the effectiveness of each approach on different types of anomalies. One of the reasons we use a synthetic data set is to be able to position the number and locations of anomalies throughout the data set. It is an idealized, but useful version of the actual data set for definitive comparison and performance evaluation. We generate 365 units of time, emulating a year’s worth of data, and artificially induce point anomalies, collective anomalies and contextual anomalies. We simulate from a Poisson distribution ($\lambda = 800$; Fricker, 2013) while randomizing the location and the magnitude of the point and collective anomalies. Figure 4.3 shows a graphical representation of the synthetic data alongside details about their associated anomalies.

To generate point anomalies, we choose 5 individual days on which to inject an anomaly. We randomly select the timing of the anomalies within the 365-day testing period using a uniform distribution. To add the anomaly, we adjust the magnitude of original, distributionally generated value by a uniformly distributed amount between –800 and 800. The magnitude is then added to the existing value at the chosen location in time.

To generate collective anomalies, we similarly randomized the location and magnitude shift applied to the original values, while also adding a randomized length of each collective anomaly from 3 to 10 days. We chose this interval based on our knowledge of existing online commercial sex work advertisement counts, which provide the basis for these data.

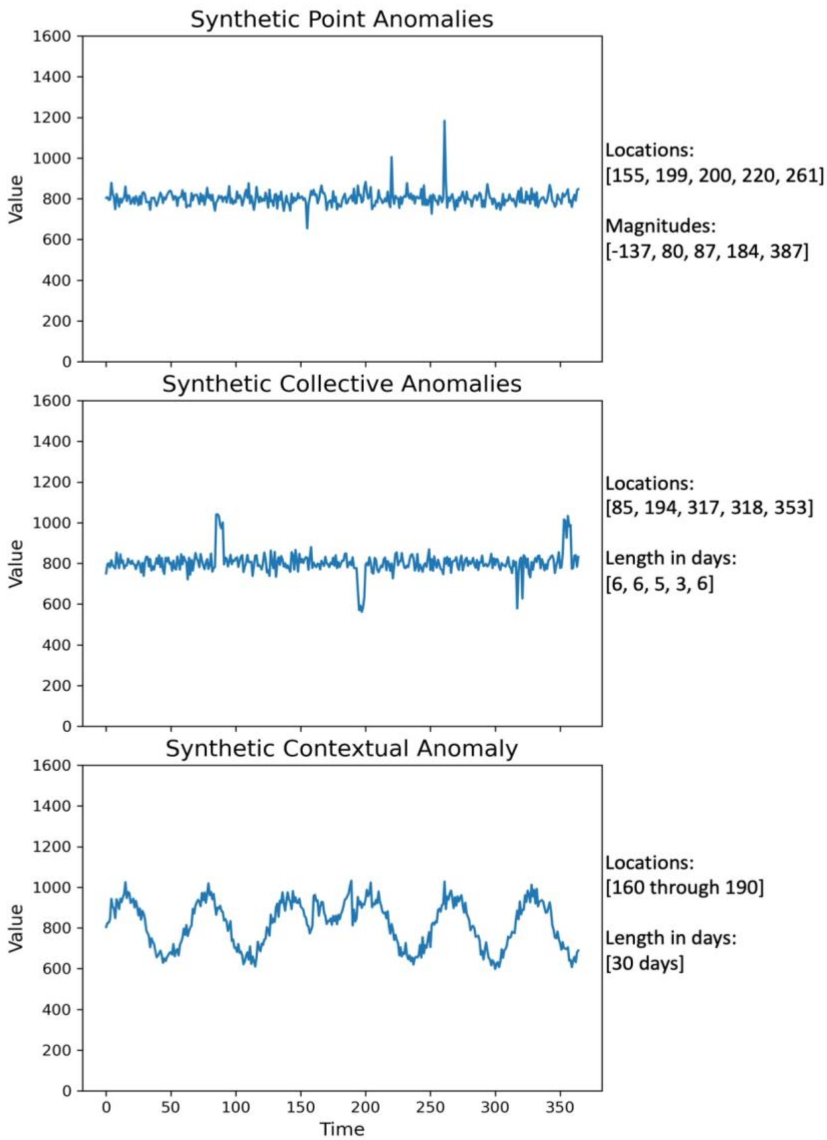

Figure 4.3 Synthetic data (left) and anomaly information (right)

The artificially generated data sets we use for analysis in Figure 4.3 follow a Poisson distribution, $p(x) = \frac{e^{-c}c^{x}}{x!}$ where x = 0, 1, 2… is the number of nonconforming occurences (anomalous data points) and $c > 0$ is the λ parameter of the Poisson distribution, which is the mean and variance. In this case our λ represents the mean and variance of online advertising counts per day.

Contextual anomalies differ from point and collective anomalies in that they require contextual changes in the data set in order for some pieces of data to stand out as anomalous. For example, electricity usage surges during the summer months as customers use more air conditioning and goes down as the temperature cools. Median electricity usage that appears within normal bounds for the entire year might be overlooked, but knowing the seasonal context of increased summer use trends, an analyst will likely identify anomalous usage. We introduced a sine curve in order to mimic seasonal changes in the data. We then artificially induced an anomalous period into the middle of the data.

4.3.2 Methods for Univariate Time-Series Anomaly Detection

We used a common approach to assess the performance of method of time-series anomaly detection being examined. Appendix Figure 4 represents an overall flowchart of all three algorithms applied to each of the three different types of synthetic data sets to arrive at the desired metrics used for evaluation. The descriptions of each algorithm in this section begins with its relevant flowchart that graphically diagrams the method for each section. These subsections culminate with the Evaluation Methods which detail how we conducted our comparison. We implemented each of the algorithms with Python using a combination of the modules Statsmodels (Seabold & Perktold, 2010), Scikit-learn (sklearn; Pedregosa et al., 2011), ruptures (Truong et al., 2020), and custom code.

4.3.2.1 Machine learning

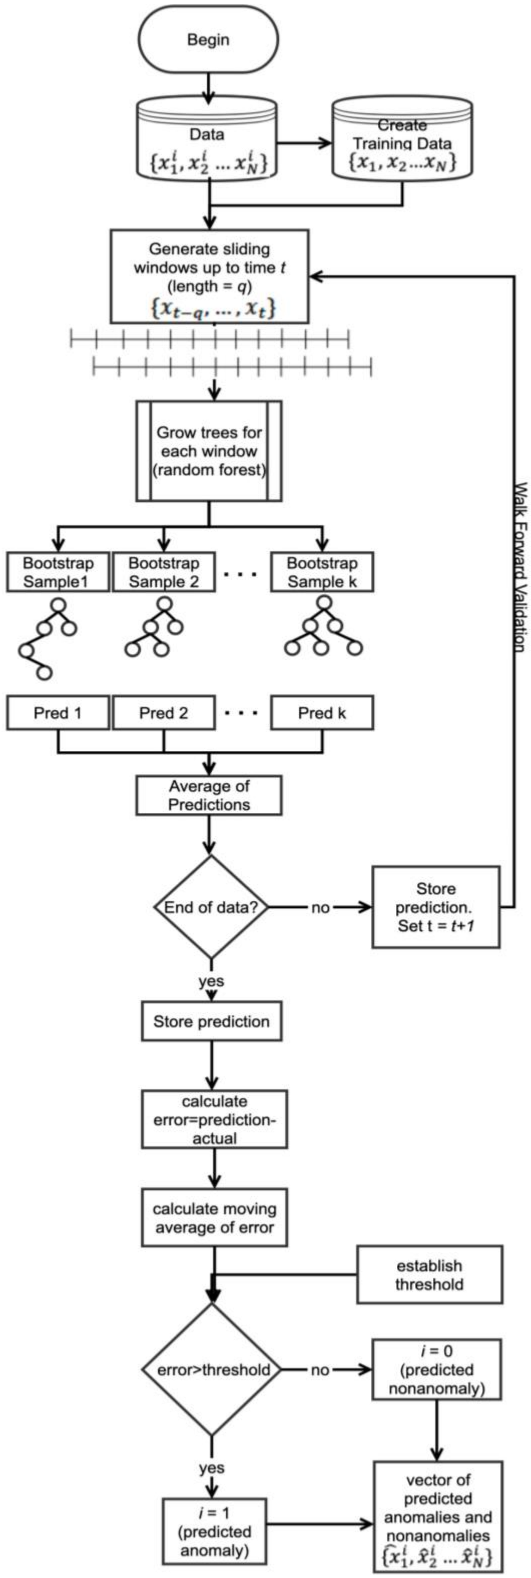

Figure 4.4 Machine learning flowchart

Input: A synthetically generated collection of time-series data with randomly generated anomalies to model system behavior. Let $X_t^i \coloneqq \{x_{t-q}, \ldots, x_N\}$ where x represent a single data point, $t =$ the current time period, q = number of one-day lags and $i \in \{1,0\}$ represents the existence of anomaly (1) or a nonanomaly (0)

Output: $\hat{X}_t^i \coloneqq \{\hat{x}_{t+1}, \ldots, \hat{x}_n\}$, the series of expected values of predictions, $t =$ the current time period, q = number of 1-day lags, n = number of days in the test set, and $i \in \{1,0\}$ represents the existence of anomaly (1) or a nonanomaly (0)

We use the synthetically generated array of values as input into a one-step-ahead prediction algorithm to produce an array of predictions (Bontempi et al., 2012; Braei & Wagner, 2020; Hill & Minsker, 2010). The random forest trains on the synthetic data set with no anomalies and predicts the first step in the test set. Anomalous data cannot be expected to produce reasonable predictions, which is why our training and testing data are split between a 365-day set of synthetically generated Poisson ($\lambda = 800$) series followed by the synthetically generated test set with embedded anomalies[2] (see Fig. 4.3). To parameterize our random forest regression model, we train our regression on 365 days of training data via a semisupervised learning method (Brownlee, 2020).[3] Using 1-day lags of length q as variables, we predict and then add the first value of the test set to the training data set and refit the model. This sliding window moves through time to predict the second value of the test set, and so on. This method, often referred to as one-step-ahead validation or walk-forward validation, preserves the temporal component of the series

[2]In both cases of the point and collective anomaly data sets, we allowed the model to train on a year's worth of synthetically generated data. For the contextual anomaly, however, we give it 2 years of data to learn the data's seasonal patterns.

[3]We characterize the process as semisupervised because we make use of labeled data that are mostly normal instead of an unlabeled data set that would be the object of unsupervised learning.

as it does not randomize the order of the data points as do other validation techniques such as k-fold cross validation (Bontempi et al., 2012; Friedman et al., 2001). Refer to Figure 4.4 for a graphical representation of the walk forward validation.

First, we use X_t historical values to train the random forest regressor. The random forest regressor uses the previous *q*-day lags as variables into the regressor. The random forest regressor depends on *t*, time, because we train a new random forest model each time with the full history up to *t* to get a "better" estimate of the parameters. The random forest ensemble method selects the parameters that minimize the prediction error across the testing set. We then use this model to predict the next period, $\hat{x}_{t+1}$; we store the predicted value, append the actual value (x_{t+1}) to the historical series, and refit the model to predict the next step. We observe the actual collection of values$\{x_{t+1}, \dots, x_N\}$ and the predicted expected values $\widehat{\{x}_{t+1}, \dots, \hat{x}_n\}$ to find if $\hat{x}$ is ever largely different from *x*. We create a new series of errors where $\varepsilon = \{(x_{t+1} - \hat{x}_{t+1}), \dots, (x_n - \hat{x}_n)\}$. If the comparison is greater than 2 standard deviations above the moving average of error, we set *i*=1; otherwise *i* = 0 ($moving\ average\ = \frac{1}{M}\sum_{i=0}^{n-1} \varepsilon_{M-i}\ where\ M = window\ size$).

Figure 4.5 graphs the actual data points from the test set with anomalies, along with the predicted values from the random forest. The simple moving average of the error is the centerline, for which we establish 2 standard deviations as our threshold to capture a reasonable number of anomalies. We tested between 1 and 3 standard deviations; 3 standard deviations benchmarks such a high standard for error that the likelihood the algorithm would detect an anomaly is low, while 1 standard deviation would be too generous in the amount of error it would tolerate and signal an anomaly too frequently (see Figure 4.5). The first two comparisons use 7 day moving average of the error for point and collective anomalies the contextual anomaly used 63 days[4].

[4] We use 63 days here because of the contextual anomaly's inherent seasonality. See also Section 4.4.1.3

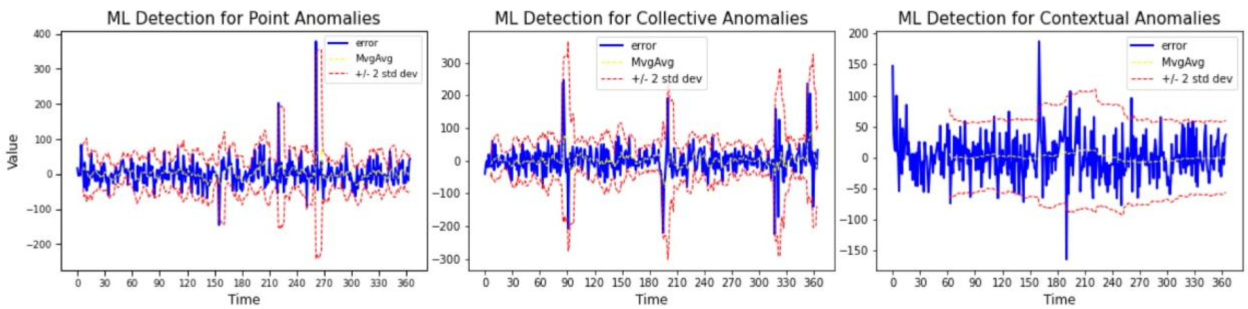

Figure 4.5: Anomaly detection for machine learning on synthetic data. Error (blue line) is the difference between actual and predicted. The red line is 2 standard deviations away from the moving average of the errors (yellow).

We use 1-day lags as variables so we are using only previous values from the same data stream. This feature of the method would be especially helpful if a risk analyst suspected gaps in their data because our model would be able to make predictions with the existing data rather than having to use complicated sampling techniques for insufficient data.

We recognize there is a vast array of techniques and learning algorithms for possible use in time-series anomaly detection. A recent survey of state-of-the-art anomaly detection techniques categorizes statistical methods like ARIMA (autoregressive integrated moving average), neural network approaches such as Convolutional Neural Networks and long short-term memory, and other classical machine-learning approaches. The authors favor classical machine learning approaches over deep-learning methods for the best balance of computational expense and accuracy (Braei & Wagner, 2020). Among the classical machine-learning approaches is the random forest. It is a powerful and common tool for classification and prediction (Omar et al., 2013). We choose random forest to represent the family of classical machine learning approaches. Random forests are characterized by an ensemble of decision trees that have one root node and many leaf nodes that use a majority vote to classify events (Breiman, 2001). They converge quickly and are robust to overfitting, making random forest a suitable learning algorithm for our walk-forward validation. (See Appendix B for more on random forest)

4.3.2.2 Change Point Detection

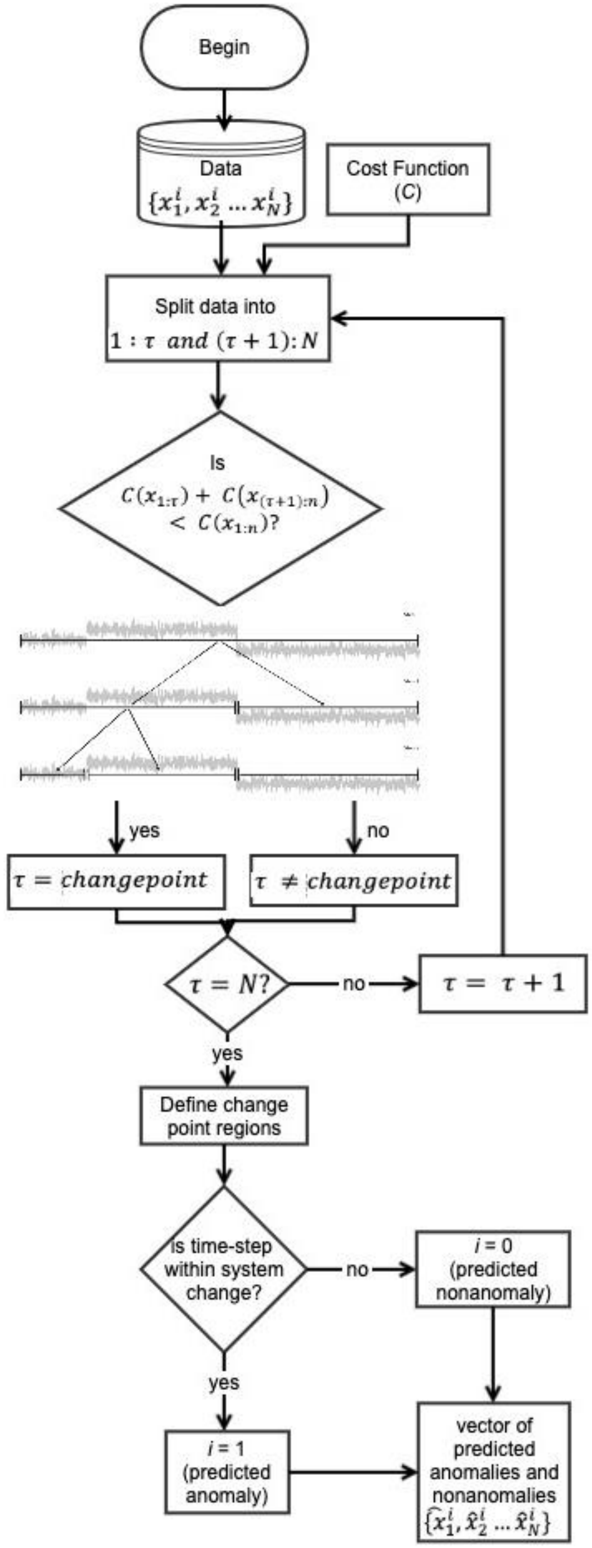

Figure 4.6 Flowchart for Change Point Detection

Input: A synthetically generated collection of time-series data with randomly generated anomalies to model system behavior. Let $X_t^i \coloneqq \{x_1^i, x_2^i \ldots, x_N\}$ where *x* represent a single data point, $t =$ the current time period, and $i \in \{1,0\}$ the existence of an anomaly (1) or a nonanomaly (0)

Output: Set of estimated changepoint indexes $\{\tau_1, \ldots, \tau_n\}$.

Change points are the locations in a time-series that mark the abrupt variation in the observed system (Aminikhanghahi & Cook, 2017). If we let $X_t^i \coloneqq \{x_1, \ldots, x_t\}$ be a sequence of time-series, then we can define change point detection as hypothesis testing between the null hypothesis H_0: no change occurs and the alternative hypothesis of H_A: change occurs where there exists a τ changepoint, $1 < \tau < t$ (Aminikhanghahi & Cook, 2017).

A vast array of change point detection algorithms exist (Killick et al., 2012; Scott & Knott, 1974; Taylor & Letham, 2018), but a recent benchmark study evaluating 14 different change point methods identified binary segmentation as receiving the highest performance measurements for univariate time-series (van den Burg & Williams, 2020). All change point detection methods can be expressed by their cost function and search method. The cost is low if the series is homogenous, with no change point, or large if the series is heterogenous, with multiple change points. Binary segmentation algorithms work to identify multiple change points by minimizing the sum of costs, $\sum_{i=1}^{n}[C(x_{(\tau_{i-1}+1):\tau_i})]$, where *C* is the cost function and $\tau_1, \ldots, \tau_{n-1}$represents each change point τ integer value between 1 and *n*-1(Rohrbeck, 2013). The algorithm starts by applying this method across the entire sequence and iterates to find a split in the series such that $C(x_{1:\tau}) + C\left(x_{(\tau+1):n}\right) < C(x_{1:n})$ where the cost function over the two subsets is smaller than the cost function of the entire series (Rohrbeck, 2013). The algorithm searches for the change point that lowers the sum of the costs. Refer to Figure 4.6 to see how we determine change points.

If such a point exists, the index is stored as a change point and the search method signals a split at that position and creates a subsequence of points, after which the same operation is repeated until no change points exist (Truong et al., 2020). In its current instantiation, the cost function we use is a least-squared deviation $\left(\sum_{t=i+1}^{n}||x_t - \bar{x}_{in}||_2^2\right)$ that compares each value with the empirical mean of the subseries ($\bar{x}_{mn}$) (Truong et al., 2020). [5]

The algorithm's output provides a location index for the first change point, τ_1. We assign an anomaly (*i=1*) for all *x* until the next change point at τ_2. We create a vector of 1s and 0s: 1s where the algorithm detects abrupt shifts in the system process and 0s where the system is behaving normally. Section 4.3.2.4, Evaluation Methods, details how we use this vector to determine the metrics used for evaluation, and Figure 4.6 provides a visualization of how the algorithm flows.

[5] For the risk-based decision-makers that cannot estimate the number of potential breakpoints in the data, a change point detection formulation that includes a bottom-up approach may be more appropriate. Bottom-up segmentation begins generously with many change points and successively deletes the less significant ones. Alternatively, the Pruned Exact Linear Time algorithm can account for an unknown number of change points by starting with a sufficiently large τ_{max}, while minimizing the sum of costs and choosing the computed segmentations that minimize a complexity penalty constraint designed to balance the goodness of fit (Truong et al., 2020).

4.3.2.3 Statistical Process Control

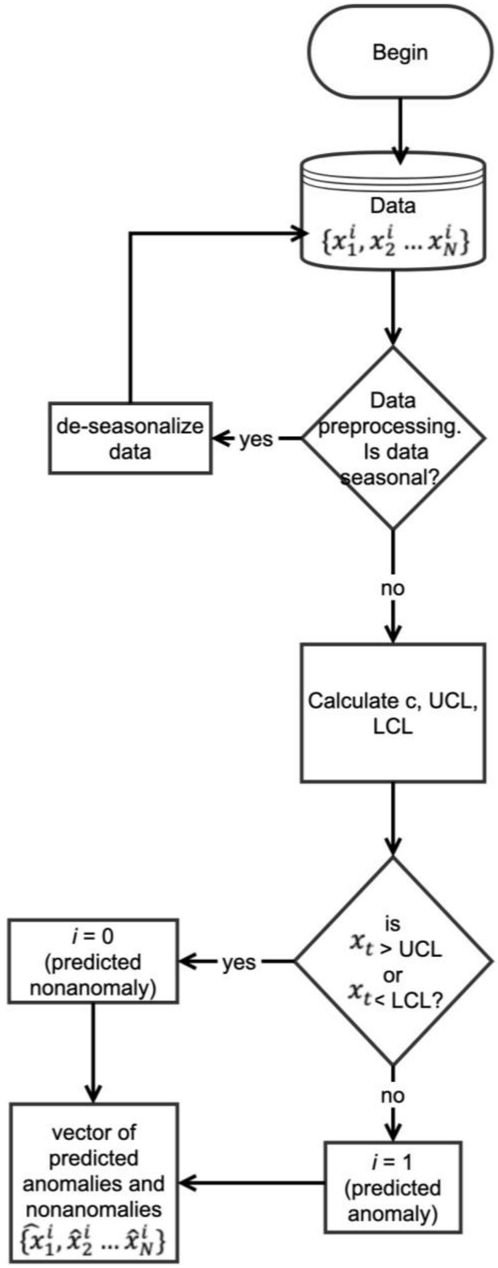

Figure 4.7 Flowchart for Statistical Process Control

Input: A synthetically generated collection of time-series data with randomly generated anomalies to model system behavior. Let $X_t^i \coloneqq \{x_1^i, x_2^i \; \dots, x_N\}$ where *x* represent a single data point, $t =$ the current time period, and $i \in \{1,0\}$ the existence of an anomaly (1) or a nonanomaly (0)

Output: Set of time indexes t where X_t^i is considered to be out of control limits.

Shewart designed control charts to provide a graphical display of quality characteristics and nonconformities generated by a process (1931). We adapt his control chart concept to apply to the artificially generated point, collective, and contextual anomalies using the same principles and formulation. Figure 4.7 shows that we implement our detection using a center line, an upper control limit (UCL), and a lower control limit (LCL). Points that exceed these control limits are evidence that the process is producing nonconforming items, is out of control, and requires investigation or corrective action (Montgomery, 2012). In order to apply Shewart's control charts, we draw parallels between this manufacturing or quality control process and our synthetically generated online commercial sex advertisements. We associate statistical process control language related to nonconformities and out-of-control processes with corresponding count-data anomalies in our series.

We use a c-chart variation of control charts because our input is discrete, with a constant sample size similar to many other applications of c-charts, such as counts of errors per document, broken rivets on an aircraft wing, or flaws on a personal computer. This control chart assumes that the data are well modeled by the Poisson distribution, in which each inspection unit represents an identical area of opportunity for a defect or anomaly to occur (Montgomery, 2012). Recall from Section 4.3.1.3 that $c > 0$ is the λ parameter of the Poisson distribution, which is the mean and variance. The definitions for our c-chart UCL and LCL are as follows (Montgomery, 2012):

$$UCL = c + 3\sqrt{c}$$

$$centerline = c$$

$$LCL = c - 3\sqrt{c}$$

If $x_t^i > UCL$ or if $x_t^i < LCL$, then i = 1, otherwise $i = 0$. This algorithm creates a set of 365 points as a vector of 1s and 0s, similar to that produced by each of the other two algorithms.

In many practical situations, the conditions necessary for a c-chart will not always be satisfied. Importantly, the contextual anomaly violates our assumption of a constant mean (see Fig. 4.3). We artificially induced a cyclical pattern into the data to simulate seasonal data and make it possible to introduce contextual anomalies. Keeping this in mind, a statistical process control approach would first attempt to subtract or eliminate the seasonal component of the data. We do not apply seasonal adjustments in data preprocessing, but rather consider seasonal adjustments on a case-by-case basis for each approach. We compare the deseasonalized data against the upper and lower control limits using the same protocol described for the point and collective anomalies. If the deseasonalized series of the contextual anomaly, point anomalies, or collective anomalies exceed the threshold, this finding will signal the need for investigation. A risk-based decision-maker can similarly use the presence of anomalies and their positions in time to signal a need for potential further examination and an action plan.

4.3.2.4 Evaluation Methods

In order to compare the performance of the machine learning, change point detection, and statistical process control techniques on our synthetic data set, we need global criteria that address the quality of each method for risk-based decision-making. A risk-aware perspective would ask what the existing risks are, how likely they are to occur, and what the consequences would be of a given unwanted outcome (Aven, 2010). Those in a position to make decisions that will mitigate risk or allocate resources to prevent risk will be concerned about whether a given method reliably identifies existing risks, and whether it falsely identifies risks that do not exist. A high-sensitivity algorithm may accurately identify anomalies but generate many false positives, resulting in low confidence that detected anomalies correspond to existing risks. We thus evaluate each method based on its accuracy, precision, sensitivity, and resulting F-score.

Numerous methods exist to evaluate anomaly detection approaches. A recent survey of the state of anomaly detection describes current evaluation criteria as including combinations of true positive rate, false positive rate, F-score, and area under the curve (Braei & Wagner, 2020). Another systematic literature review of metrics for anomaly detection and prediction finds that commonly used criteria include distance indexes, histogram visualization, precision, recall, F-measures, positive predictive value, and mean absolute error, among many others (Fahim & Sillitti, 2019). Some biosurveillance methods advocate for measures such as average time between false signals, probability of successful detection, and condition expected delay to observe outbreaks for early event detection (Fricker, 2013). Figure 4.8 shows how we use a combination of metrics drawn from the literature that we can universally apply to all three algorithms.

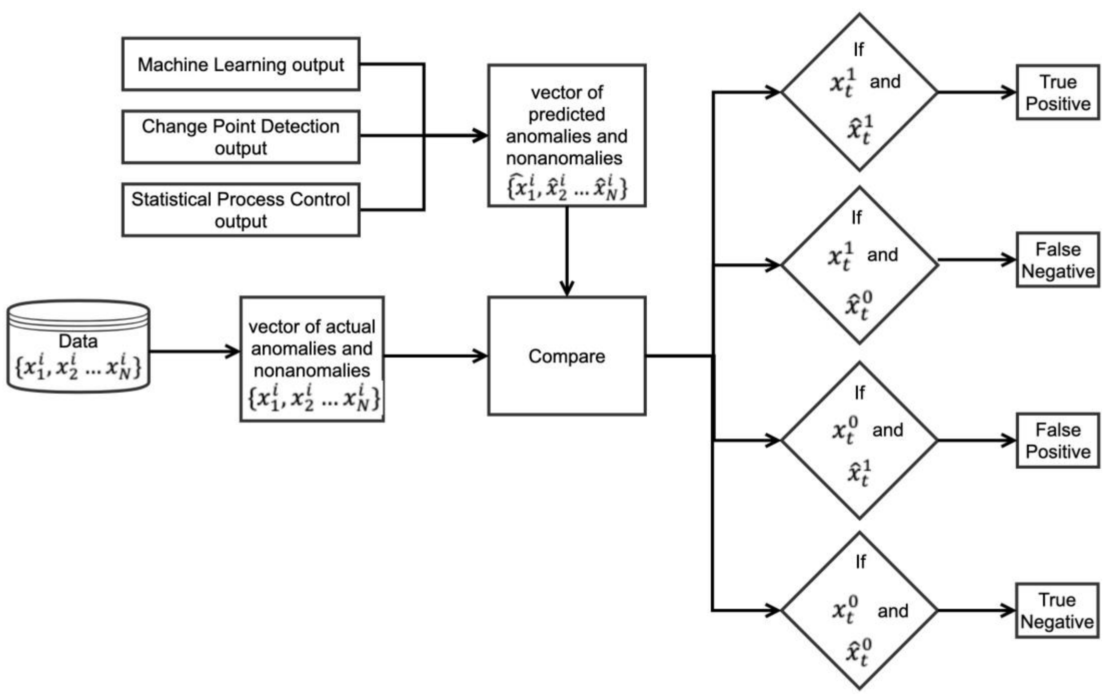

Figure 4.8 Flowchart of generalized algorithm comparison

This research uses counts of Type I and Type II errors to build metrics that will evaluate the different time-series anomaly detection approaches (see Table 4.1). Each algorithm produces a vector of hypothesized anomalies indicated by a simple yes/no (1 or 0) decision. This vector is then compared with another vector of the "true" anomalies from the synthetic data set[6]. We draw results for both a daily comparison metric and a window comparison metric. The daily comparison metric focuses on a one-to-one comparison for each individual day, whereas the window comparison metric uses an incremental window across the whole data set. We use both in our analysis because the daily comparison metrics boost our accuracy measurement because of our imbalanced data set. Anomalies are rare, so we unsurprisingly see a large number of true negatives across all approaches (where the algorithm correctly identifies the datum as a nonanomaly).

The window comparison metrics account for this imbalance by clustering the points. Along the same lines, the daily comparison metric penalizes algorithms that do not identify the anomaly on the precise day, but lag a few days in the window. The window comparison metric accounts for this penalty by being slightly more generous regarding when the algorithm detects the anomaly. For example, if the predicted vector identifies an anomaly within a 5-day increment of the identified anomaly in the actual vector, then the entire window is classified as a true positive. We used this adaptation of a time window to provide a margin of error and improve accuracy of the algorithms suggested in literature (Farshchi et al., 2018; Truong et al., 2020). Figure 4.9 is a simplified visualization (not to scale) of this counting protocol for this window comparison.

[6] The only exception to this protocol is for the contextual anomaly using statistical process control. As a result of the seasonality of that data set, we evaluate the transformed data instead of the raw values.

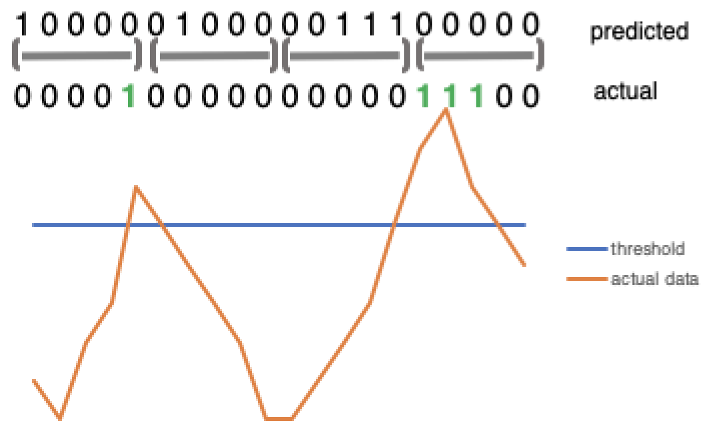

Figure 4.9: Visualization of window comparison. A vector of 1s and 0s representing the actual anomalies and the predicted/labeled anomalies allows them to be compared. If the actual anomaly exceeds the threshold it is labeled as a 1. If the model predicts an anomaly in the predicted vector window, the window is labeled as a 1.

A point (or collection of points) in our synthetic data set that the algorithm labels as an anomaly, but is not an anomaly, is counted as a false positive. In Table 4.1 we offer our "confusion matrix," which presents the definitions that drive our ratios and evaluation, adapted from the literature (Aminikhanghahi & Cook, 2017; Palaniappan et al., 2012).

Table 4.1: Confusion matrix used anomaly detection classification

	Classified as anomaly	Classified as nonanomaly
True Anomaly	TP (True-positive)	FN (False-negative)
True Non-anomaly	FP (False-positive)	TN (True-negative)

In order to provide a high-level impression of each algorithm, we define accuracy as the ratio of the correctly identified data points to the total number of data points.

$$Accuracy = \frac{TP + TN}{TP + FP + FN + TN}$$

While the accuracy rating is based on the proportion of correct predictions versus total predictions, the precision rating is based on the proportion of correct positive predictions (true positives) versus all points classified as an anomaly.

$$Precision = \frac{TP}{TP + FP}$$

Sensitivity, also described as Recall or True Positive Rate, evaluates the proportion of true positive anomalies versus total anomalies (including those not identified as anomalies by the system).

$$Sensitivity = Recall = \frac{TP}{TP + FN}$$

Finally, the F-score provides a broad measure of the algorithm performance and indicates if the method produces low false positive and low false negative outcomes. A perfect F-score would be 1, whereas a complete failure would be 0. This metric is a weighted average of the true positive rate (recall) and precision.

$$F\ Score = 2 * \frac{Precision * Recall}{Precision + Recall}$$

The purpose of using a combination of all of these measurements is to provide balance when comparing the different algorithms. In general, accuracy can be a sufficient overall measurement of the algorithm’s performance, but in an imbalanced data set, we calculate supplemental metrics to help characterize performance. In order to improve precision, we could lower the threshold, allowing a given method to detect more anomalies, but this improvement would come at the expense of accuracy and sensitivity. It is thus important to take all these measures into account.

4.3.2.5 Limitations of Methods

Despite a thorough examination of three potential algorithms for use in risk-based decision-making, we acknowledge that our method has limitations. We categorize these into two areas: 1) interpretability and 2) data. The most salient limitation is in the interpretability of metrics. We will observe in the Results section that some algorithms are better than others at identifying certain types of anomalies. This disparity would result in an analyst receiving conflicting results if they used multiple methods. In another context, a researcher could take the mean across all different algorithms, but we rely on binary classification, making that approach unrealistic. One could, however, take a majority-rules approach and address only anomalies identified by most of the algorithms. None of these methods is perfect. In short, the interpretability of the metrics lies in the ability of the risk-based decision-maker to implement algorithms in a way that balances the relative cost of missing an anomaly against the cost of investigating one.

The other notable limitation in this algorithm comparison is the requirements for the data used for examination. As discussed in the Background section, we contextualized our algorithms in an offline setting, examining the data in batch retrospectively. If we were to apply these methods online and allow the real-time incorporation of new data streaming into the system, our algorithms would change significantly. Such a configuration could allow a timelier response, producing a prediction before the next time step occurs. Our data are also constrained by requirements related to dimensionality and continuity. Our synthetic and real-world data sets are univariate series for use in our comparison, but risk analysts may find themselves in a position of having to analyze multivariate data. In most cases, the proposed algorithms can be modified to accommodate multiple variables, and many studies exist in the literature that address multivariable anomaly detection (Boecking et al., 2019; Dubrawski, 2011; Fricker, 2013; Goldstein & Uchida, 2016; Matteson & James, 2014). For example, Boecking et al.'s paper on online escort advertising uses a multivariable approach. It analyzes multiple streams of data and detects an anomaly by using a contingency table to compare a baseline window with a reference window over the time-series

(Boecking et al., 2019). Lastly, missing data and sampling errors can occur in the collection of data. We found this to be true in the scraping efforts used to obtain our real-world data set. Similarly, risk analysts attempting to use anomaly detection techniques would need to determine if their analytic sample is missing data and whether or not it is still representative of the larger population. They may need to incorporate missing data techniques such as imputation, deletion, maximum likelihood estimation, or simulation in the preprocessing step to mitigate this limitation (Enders, 2010).

4.4 Results

This section compares the three different approaches we implemented for time-series anomaly detection.

4.4.1 Comparative Evaluation

We begin with a visualization of the results of the algorithm comparison. All three types of anomalies (point, collective, and contextual) were evaluated using three different approaches: machine learning, change point detection, and statistical process control. Further details on each are discussed within this section.

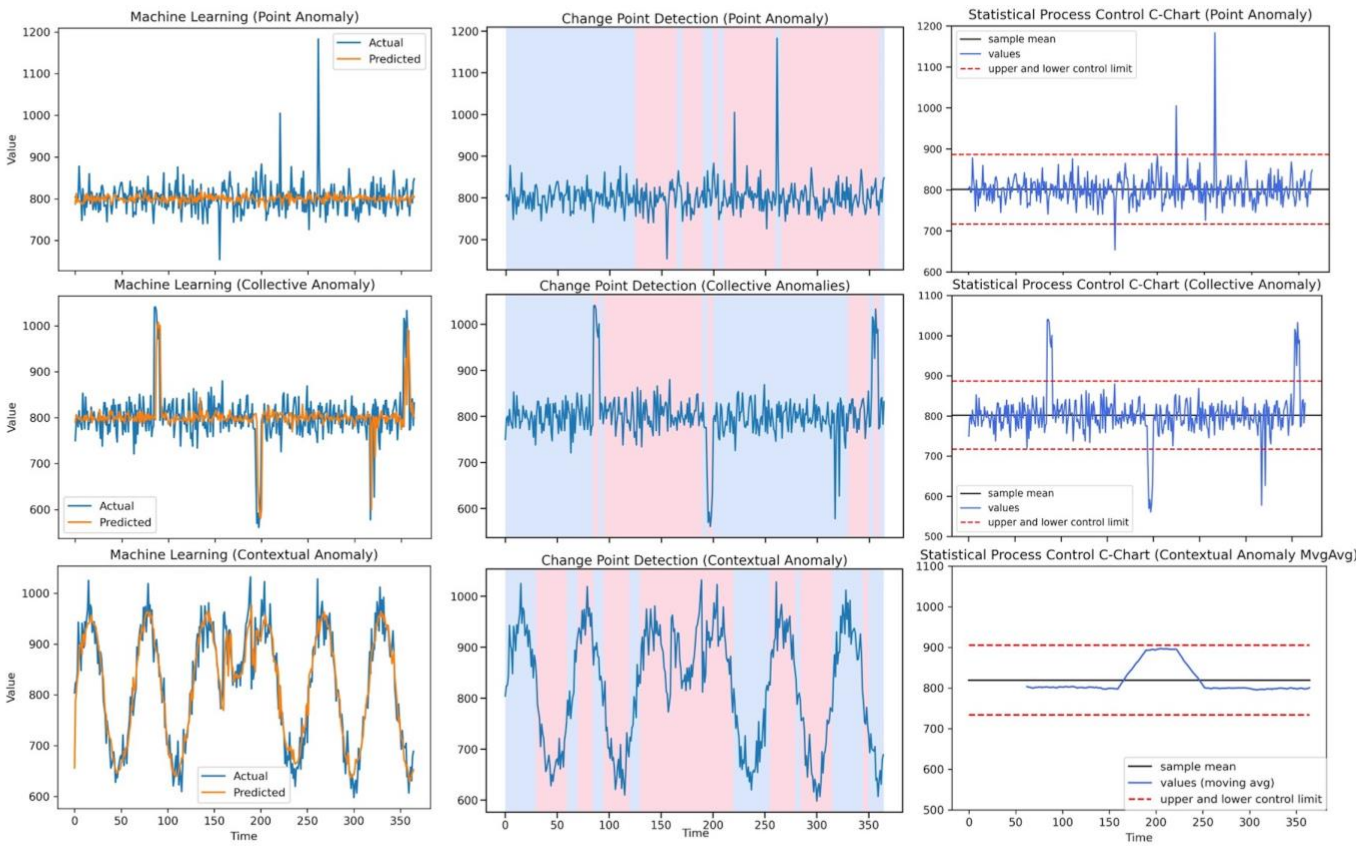

Figure 4.10: Visualization of algorithm comparison results.
Columns represent machine learning, changepoint detection, and statistical process control approaches on point, collective, and contextual anomalies (rows)

4.4.1.1 Machine Learning Results

As discussed, the machine learning experimental setup uses a synthetic training data set, and then generates predictions using a random forest on a test set. The machine learning approach differs from the other two methods in that it produced a forecast of predicted values. This differentiation also affects our analysis, in that the prediction's mathematical distance from the actual values is what allows us to determine if there is an anomaly. The first column in Figure 4.10 graphically depicts the results of the machine learning predictions (orange) along with its actual (blue) test data.

Table 4.2 is a summary of the metrics gathered for the machine-learning method, and is representative of the subsequent output of our analysis for each method type. Recall that those points that exceed the 2 standard deviation threshold over our moving average are labeled with a 1, and others are 0. We compare this with our vector of true anomalies and observe the number of true positive, false negative, false positive, and true negative data points for both the daily and window comparison metrics[7].

Table 4.2: Summary of the accumulated metrics and detection outcomes for the machine learning method on all three types of synthetic data. Anomalous "days" are out of 365 units. Daily and window refer to the daily comparison metrics and window comparison metrics

Machine Learning	Point		Collective		Contextual	
Anomalous "days"	5		23		30	
	daily	window	daily	window	daily	window
True Positive	3	3	4	4	2	2
False Negative	2	2	19	2	28	1
False Positive	3	3	2	2	12	10
True Negative	357	29	340	29	323	18
Accuracy	0.986	0.865	0.942	0.892	0.89	0.645
Precision	0.5	0.5	0.667	0.667	0.143	0.167
Sensitivity	0.6	0.6	0.174	0.667	0.067	0.667
F-Score	0.545	0.545	0.276	0.667	0.091	0.267

Accuracy and sensitivity scores are both important in interpreting these results. Risk analysts and policy-makers will want to know if the algorithm detects a true anomaly accurately, but will also want to avoid a false negative, in which the algorithm identifies the point as a nonanomaly where it is indeed an anomaly. This is best represented by the sensitivity measurement. We can be slightly more generous with false positives, as they trigger an alarm that

[7] The window comparison metrics all use a 10-day window with the exception of the contextual anomaly, which uses a 12-day window. We tuned parameters to give each method the best chance at finding anomalies, and favored window lengths that found more anomalies. When window parameters did not make a significant difference we chose 10 days.

would initiate an investigation, but the risk analyst would need to balance the cost of that investigation with the cost of missing an anomaly. As expected, we see a reduction in accuracy using the window-based comparison metrics because we have artificially reduced the number of true negatives by clustering them together in 10-day increments.

This made it possible to obtain a more comprehensive understanding of the method's performance. The dramatic difference in sensitivity scores for the collective anomaly comparison (0.174 and 0.667) is likely due to the fitness of the model. As discussed in the Methods section, the collective anomalies range between 3 and 10 days. The model learns after the first spiking day that the leading edge of this anomaly will be followed by a short series of similar points, thereby reducing the error that these metrics are measuring. This also demonstrates the usefulness of the window comparison metric, which produces a higher and perhaps more accurate sensitivity score for the collective anomaly (0.667). The window comparison metric gives credit to the entire window after finding the first spike, which in a real-world scenario would be helpful in identifying a period of time that merits investigation. Lastly, the metrics for the contextual anomaly also benefit from the fitness of the model, in that here we see it captures the beginning and the end of the contextual anomaly but does not identify the intervening period as an anomaly.

4.4.1.2 Change Point Detection Results

The graphical output for the change point detection method is summarized in Figure 4.10. The contrasting blue and pink blocks indicate the associated nonanomalous (blue) and anomalous (pink) periods determined by the binary-segmentation algorithm. Table 4.3 below summarizes the collected metrics.

Table 4.3: Summary of the accumulated metrics and detection outcomes for the change point detection method on all three types of synthetic data. Anomalous "days" are out of 365 units. Daily and window refer to the daily comparison metrics and window comparison metrics

Change Point Detection	Point		Collective		Contextual	
Anomalous "days"	5		23		30	
	daily	window	daily	window	daily	window
True Positive	3	4	14	4	30	3
False Negative	2	1	9	2	0	0
False Positive	207	19	116	11	125	13
True Negative	153	13	226	20	210	21
Accuracy	0.427	0.459	0.658	0.649	0.658	0.649
Precision	0.014	0.174	0.108	0.267	0.194	0.188
Sensitivity	0.6	0.8	0.609	0.667	1	1
F-Score	0.028	0.286	0.183	0.381	0.324	0.316

The change point detection method performed similarly to the machine learning approach at detecting true positive point anomalies; this is reflected in the sensitivity score. Overall accuracy, precision, and F-score measurements, however, are much lower due to the high number of false positive labels this method generates for the data set. For collective anomalies, the change point detection method performs relatively well, finding more than half of the collectively anomalous windows in both the daily and window comparison methods. Again, the presence of many false positives interferes with the precision and F-score for all three different types of anomalies. Similarly, although the change point detection method identifies many false positives compared to the contextual anomaly also, it does very well at identifying all of the contextual anomalies, resulting in a perfect sensitivity score. All of these window comparison metrics used a 10-day incremental window.

The data set reflecting the contextual anomaly is particularly useful in illustrating the balance between oversegmentation and undersegmentation (Truong et al., 2020). In Figure 4.10 we see that the change point detection algorithm's implementation on the contextual data appears to capture all of the cycles instead of just the anomalously noncyclical period beginning at day 160. The incorporation of more change points into the algorithm will allow analysts to capture the anomalous period, but at the expense of precision and the addition of more segmented periods that are ultimately labeled as false positives.

4.4.1.3 Statistical Process Control Results

Figure 4.10 displays the graphical output of the statistical process control c-chart used for each of the series associated with point, collective, and contextual anomalies. The dotted red lines indicate the upper and lower control limits calculated from the formulation from the Materials and Methods section. The solid black line indicates the centerline, or c. The solid blue lines for the point and collective anomalies under the statistical process control method indicate the synthetic data set in its raw form. In the contextual anomaly (lower right), however, the solid blue line indicates the moving average of the data.

Recall that detecting the contextual anomaly requires understanding seasonality due to the datas' cyclical behavior. In order to deseasonalize the data, we took a moving average of the length of the cycle, which is 63 days. If the data had no anomalies, we would expect a relatively flat line after removing the cycles. Given that we have a known anomaly at 160-190 days, we hypothesize a rise in the data slightly lagging after the known anomalous period. This is precisely the behavior we witness in Figure 4.10's relatively flat deseasonalized line with an a suspected anomalous period associated with the 30 day anomalous period. Although we recognize the pattern in the data may be worth investigating, the deseasonalized data still do not exceed the upper control limit threshold. This is also clear in Table 4.4, which summarizes the detection metrics.

Table 4.4: Summary of the accumulated metrics and detection outcomes for the statistical process control method on all three types of synthetic data. Anomalous "days" are out of 365 units. Daily and window refer to the daily comparison metrics and window comparison metrics. (indicates undefined, dividing by 0)*

Statistical Process Control	Point		Collective		Contextual	
Anomalous "days"	5		23		30	
	daily	window	daily	window	daily	window
True Positive	3	3	20	6	0	0
False Negative	2	2	3	0	30	3
False Positive	1	0	0	0	0	0
True Negative	359	32	342	31	335	34
Accuracy	0.992	0.946	0.992	1	0.9178	0.9189
Precision	0.75	1	1	1	*	*
Sensitivity	0.6	0.6	0.87	1	0	0
F-Score	0.667	0.75	0.93	1	*	*

In regard to the point anomalies, the statistical process control method performed relatively well compared to the machine learning method and better than the change point detection method because it more accurately labeled true negatives and false positives. This method was particularly useful in correctly labeling all of the collective anomalies. Although in the daily comparison metrics it was able to find almost all of the global anomalies, it was less effective at finding the local anomalies. This is due to the nature of established thresholds using the control limits. On the other hand, using the window metrics on the collective anomalies with the c-chart, we are able to correctly identify all of the existing anomalies.

4.4.2 Performance Comparison and Discussion

The analysis did not identify one clear dominant approach for all forms of univariate time-series anomaly detection for risk-based decision-makers. Instead, we found that each method's detection power varied with the different types of anomalies. The statistical process control method using a c-chart performed well for both the point and collective anomalies and very poorly for the contextual anomalies. It is the most powerful approach among the three algorithms at finding collective anomalies when combined with the window comparison method. When the upper and lower control limit threshold for anomalies is set correctly, a risk analyst can be confident that this approach is relatively effective at discovering anomalies that are grouped together collectively. Importantly, this method used for collective anomalies scored very high in precision as compared to other methods that labeled more false positives. The likelihood of false positives is an important consideration for a risk-based decision-maker who uses these metrics to appropriate resources. If investigating outcomes is a costly and time-consuming task, one would want to limit the number of false positives and choose a method higher in precision.

The metrics for change point detection, on the other hand, suffered from a great deal of false positive labeling. Despite this shortcoming, it is the only method that evidenced robust findings for contextual anomalies, as evidenced by its ability to capture the entire length of the datas' 30-day anomaly. It does so, however, while identifying many false positive signals, which may lead a risk-based decision-maker to spend scarce resources investigating those alerts. The high degree of false positives may be due to the cyclical nature of the contextual data set and the nature of the binary segmentation algorithm's inclination to detect a distributional structural change. When a major shift occurs, it will capture all those points in that hypothesized shift. This may also be exacerbated by the use of the window comparison method, which may also penalize an entire increment. Although change point detection performed strongly in detecting a contextual anomaly, it is important to note that that contextual anomaly would still require a human analyst with domain expertise to interpret the anomalous behavior of the data. This is true of any of the algorithms being presented, but is important to consider in contexts where a series of data points might be considered normal until one understood the anomalous context.

The machine learning approach consistently performed the best at discovering all different types of anomalies among all three different methods, with the highest mean accuracy score. Although it scored higher than the change point detection method at identifying contextual anomalies, it identified only the beginning and end of the anomalous period and failed to label the intervening period. The random forest forecast learns the patterns of the data, which helps a risk analyst to identify issues at certain points in time, but it may fail to give information on the magnitude of the issue if the following points are labeled as nonanomalous. The strength of the machine learning approach is in its ability to use 1-day lags as variables, and generally handle a more complex data set seamlessly than the other two methods in the way we have presented them. In all three types of anomalies, however, without data-driven covariates to drive the understanding and training of the model, it is difficult to assess the meaningfulness of these anomalies without expert human judgement.

4.5 Online Commercial Sex Work Advertisements as a Case Study

4.5.1 Introducing Sex Trafficking as a Risk Analysis Problem

As we discussed above in the Results section, a risk analyst would approach a complex risk-laden problem by investigating three distinct areas: the risks of concern, the likelihood of those risks, and their consequences (Aven, 2010). Recall from Section 4.2.2 that in many risk analysis fields, data conform to the following features: 1) they are temporal, 2) they are univariate, and 3) the risky events are relatively rare. Our real-world sex trafficking data set demonstrates these qualities, making it conducive for a useful case study.

Sex trafficking is a global crisis that accounts for approximately a third of the global problem of human trafficking (ILO, 2017), which in the United States it is defined as the sexual exploitation of another person through “force, fraud or coercion” or including a minor (TVPA, 2000). The health risks of sex trafficking are severe, especially to its victims. Risks to communities and security risks are also well documented in the literature (Anthony 2018; Bigio & Vogelstein, 2019b; Kara, 2017; Lederer & Wetzel, 2014; Shelley, 2012). In this case study we aim to critically analyze a potential sex trafficking anomaly to illustrate the use of these algorithms against a real-world problem and to provide confidence in policy-makers’ decisions to allocate resources to the problem. Among many factors, an anti-sex-trafficking risk-based decision-maker can process the known risks coupled with the anomaly detection outputs in order to formulate an action plan to combat the crime. This type of evidence-based situational awareness could drive change in many ways including reinforced social services, policing strategies, entitlements for vulnerable populations, policy reform, or legislation. No matter the outcome, working towards solutions that provide quantitative rigor and risk-based decision-making to reduce the prevalence of sex trafficking is a worthwhile endeavor.

In addition to the previously discussed anomaly detection that occurs in the biosurveillance and manufacturing domains, social sciences have embraced the potential of detecting anomalous behavior as it relates to social networks and social media analysis (Benigni, 2017; Khamparia et al., 2020; Savage et al., 2014). Those studying social issues such as human trafficking, however, have rarely harnessed the advantages of these quantitative methods (Gozdiak, 2011). In the United States, commercial sex is often advertised online on paid websites where commercial sex workers offer their in-person services. These websites provide buyers a way to connect by phone or email; we use that information as a source for data-driven analysis. The density of advertising on these types of websites can serve as a proxy for the amount of supply and demand in a given region (Boecking et al., 2019). Mainstream media often cite major social disruptions such as the Super Bowl as a mechanism to drive supply and demand of commercial sex, and in turn, sex trafficking (Martin & Hill, 2019). The COVID-19 pandemic is another social disruption that exemplifies how anomaly detection can influence decision-makers. Media outlets reported that the pandemic's social distancing restrictions were having a negative impact on both commercial sex workers and sex trafficking victims (Okeowo, 2020). We hypothesize that established methods for anomaly detection can identify a decrease in online commercial sex advertising associated with sex trafficking following the COVID-19 stay-at-home orders of March 2020. A scientific perspective on the existence of an anomaly could make it impossible for a decision-maker to ignore the severe and disproportionate impact a pandemic could have on sex trafficking victims.

4.5.2 Anomaly Detection in Online Advertisements

4.5.2.1 Real-World Data Set

We scraped data from the website Rubratings.com in the Python programming language with the use of the BeautifulSoup package. This facilitated a regular automated scrape of the website from 140 cities across 47 states in the United States from July 2019 to May 2020. The online platforms used to sell commercial sex activities are notoriously unstructured, and often actively blocking scrapes. Mapping an unstructured website with no application programming interface made for challenging data collection and did not allow us to capture all advertisements (Boecking et al., 2019). This analysis filters out the periods during which the scraper was not actively collecting data and recognizes that the missing data could have an impact on the analysis.

The mere presence of an advertisement on this website does not indicate that the sex worker whose services are being advertised is trafficked. We employ the use of a data stratification protocol detailed in Appendix B that characterizes an advertisement as originating either from an independent commercial sex worker or a sex trafficker. Our procedure uses phone number clustering on the website to flag an advertisement as possibly related to sex trafficking. After removing duplications and gaps where the website blocked the scraper or failed to run, 235,841 total suspected sex trafficking advertisements remained across the United States for the 221-day period covered. We further aggregated this to a daily count consistent with a univariate time-series.

4.5.2.2 Analytic Approach Using Three Different Anomaly Detection Algorithms

The machine learning, change point detection, and statistical process control methods applied here are identical to those described in the Materials and Methods descriptions, with a few exceptions. As discussed, time-series data are different from other types of labeled data that are common among machine learning methods. Univariate time-series data are collected on the same metric observed over regular time intervals. This distinction is important because of the nature of the suspected sex trafficking advertising data, which conforms to a count of advertisements across the United States aggregated by day. As a result of these features, the data must pass checks for stationarity for our machine learning algorithm. This algorithm is the only method of the three that uses past data to predict future data, but we can trust such predictions only if we assume that the data will follow the same general trends and patterns as in the past. We check for stationarity in this case through a series of qualitative and quantitative checks that are detailed in the Materials and Methods section. After determining that the data are not stationary, we perform data transformations to correct for nonstationarity and implement our anomaly detection through machine learning forecasts and post analysis. This process requires an additional step of back-transformation prior to collecting metrics.

In the case of our statistical process control algorithm, if no standard is given for c, *centerline,* as is the case with our real-world data, then the observed average number of nonconformities in a preliminary sample can be used (Montgomery, 2012). In this case we use $\bar{c}$, the mean of counts in the hypothesized nonanomalous subseries, to replace c in the UCL and LCL formulations. Recall that we also applied a moving average to the synthetically generated contextual anomaly to account for seasonality. We observed in our daily counts a subtle weekly cycle, with the density of online advertisements peaking into the weekend and dipping on Sundays. For this reason, we apply a 7-day moving average to the data prior to inputting them into the control chart. We hypothesized that there exists at least one subseries after the implementation of the COVID-19 March 2020 stay-at-home orders where we would expect the algorithms to detect a dramatic decrease in online advertising (Kaiser Family Foundation, 2020). We test if each algorithm discovers anomalies starting from March 20, 2020, which equates to point 149 along the x-axis of time. (March 19, 2020 = 1219 advertising counts; March 20, 2020 = 249 advertising counts). Furthermore, we acknowledge at least two scraping anomalies (see Appendix B) that may alert the detection methods to signal a change in the system or an anomaly.

4.5.3 Results and Discussion

After running the machine learning algorithm, change point detection, and statistical process control algorithms we generated three different visualizations, shown in Figure 4.10.

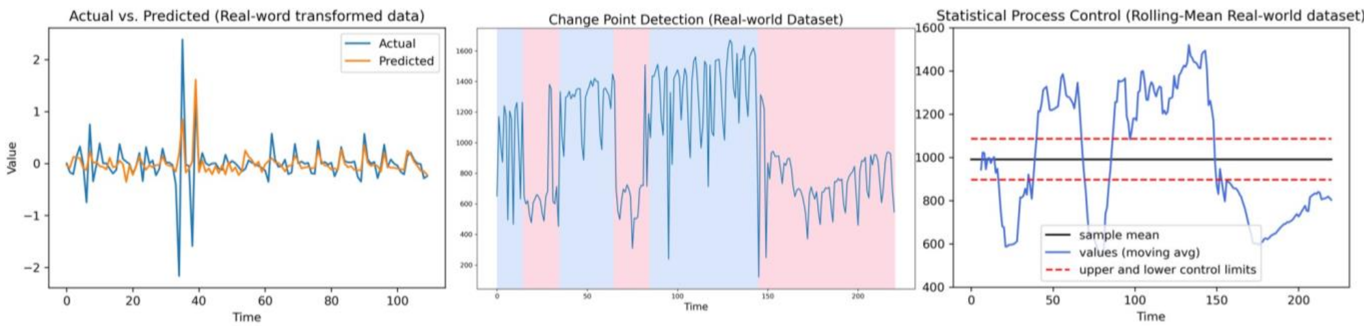

Figure 4.11: Anomaly detection for real-world data set with all three methods

In order to summarize the output of the analysis, Table 4.4 displays the collected metrics associated with running the three algorithms on the real-world online commercial sex work advertisement counts that also reflect our sex trafficking indicator. These metrics reflect both the daily and window comparison method detailed earlier.

Table 4.5: Summary of the accumulated metrics and detection outcomes for all three algorithms with the real-world data set. Anomalous "days" are out of 365 units. Daily and window refer to the daily comparison metrics and window comparison metrics. (indicates undefined, dividing by 0.)*

Real-world Dataset	Machine Learning		Change Point Detection		Statistical Process Ctrl	
Hypothesized anomalies	70 of 110		71 of 221		71 of 221	
	daily	window	daily	window	daily	window
True Positive	0	0	71	15	70	15
False Negative	70	15	0	0	1	0
False Positive	2	2	45	9	124	27
True Negative	38	5	105	21	26	3
Accuracy	0.345	0.459	0.796	0.800	0.434	0.400
Precision	0.000	0.000	0.612	0.625	0.361	0.357
Sensitivity	0.000	0.000	1.000	1.000	0.986	1.000
F-Score	*	*	0.759	0.769	0.528	0.526

From the point of view of a risk-based decision-maker, the best algorithm is change point detection, while the statistical process control approach delivered intermediate results and the machine learning approach was least helpful in drawing a conclusion. The change point detection method's performance was superior at detecting anomalies in the suspected sex trafficking advertisements. Although the statistical process control approach discovered most of the true positive anomalies, which boosts the sensitivity score, it also reported many false positives. We can also see from Figure 4.10 that the change point detection algorithm discovered two additional abrupt shifts in mean. These two shorter periods are consistent with scraping blocks and failures that would have pulled daily advertising counts down.

It stands to reason that the c-chart would not perform well in this instance. In practice, although it is not always possible to meet the condition of a constant mean for use in a c-chart, we did attempt to smooth out the data using a 7-day moving average. Figure 4.10 clearly shows that most of these points would be considered out of control. Additionally, the machine learning algorithm did not detect any hypothesized anomalies. Recall that the metrics generated from this algorithm compare the error between actual and predicted values. The failure to detect any hypothesized anomalies could be due to the lack of evidence that one exists, or a result of overfitting the data, resulting in low errors. Considering the labor-intensive preprocessing and postprocessing steps of using the machine learning approach against this type of real-world data, it may not best serve a risk analyst.

Given this evidence, a risk-based researcher can draw a few conclusions. First, this particular data set drawn using our scraping procedure may be best suited for use in conjunction with change point detection. Second, the identified anomalous periods included two scraping mishaps and the abrupt shift in suspected sex work advertising associated with the COVID-19 stay-at-home orders. These detections would lead a risk professional to further investigate these locations in time, and they would be correct to do so. The alert of this anomaly should signal a risk-based decision-maker to unpack the cause of the hypothesized anomalously decreasing period of suspected sex trafficking. At a minimum it should prompt the analyst to dig deeper with questions such as, Is this decrease a yearly issue? Do other factors exist that can account for a dramatic decrease?

Recall that detecting this type of anomaly typically requires analysts with domain expertise to play a role in postprocess analysis to interpret the algorithm's findings. Assigning less importance to the two scraping failures and more importance to the sharp decrease following stay-at-home orders is a good example of such insight. Although it is not a comprehensive analog to the synthetic data, our real-world data set for suspected sex trafficking advertisements includes a shift in data best described as a contextual anomaly. It stands to reason that the method that performs best for the synthetically generated contextual anomaly would also perform best with our real-world sex trafficking data.

4.6 Conclusion

Anomalous behavior recognition can reduce risks, identify unseen problems, and help anticipate future risks (Fahim & Sillitti, 2019). Univariate time-series anomaly detection serves as a valuable form of analysis for risk-based decision-makers. Borrowing from the biosurveillance, manufacturing, and quality control domains, we offer three different algorithms for potential use for complex risk-based decisions. This work is by no means a comprehensive analysis and cannot offer a mature management decision support system or tool, but does provide a reliable comparison that takes advantage of the well-researched disciplines of machine learning, change point detection, and statistical process control.

Our experiments using synthetically generated point, collective, and contextual anomalies have some notable findings: 1) machine learning performed the most consistently across all different types of data sets; 2) contextual anomalies were best detected by the change point detection algorithm; 3) statistical process control was superior to the other methods at discovering the collective-type anomalies. Real-world data are messy, so we sought to apply these methods to a real data set involving the complex risk-based problem of sex trafficking. We observed that in the uncertain and noisy backdrop of the real world, the machine learning and statistical process control algorithms were not as effective as the change point detection method.

Future research should incorporate additional sophisticated learning algorithms and techniques that can process diverse new streams of data for analysis. If risk-based decision-makers and risk researchers are open to using these time-series anomaly detection methods for the various risk-related challenges in our future, they will first have to access relevant and meaningful data. The complexities of the data, unexpected noise, and the imbalanced nature of rare events will require deep examination to achieve insights that can drive resource allocation and decision-making. Well-designed methods will provide a basis for well-curated data sets to derive value and knowledge. This work can go a long way in helping to provide that foundation for complex, inherently risky systems.

Chapter 5

Conclusion

5.1 Summary of Contributions

This body of research attempts to derive meaning and insights from hidden data—data that were never intended to be analyzed in this way. Doing so required the development of methods for data collection, data handling, and transforming very unstructured data into structured data. This was done in service of obtaining a deep understanding of information that has been purposefully obscured. Though advertisements for sexual services are designed to conceal the possibility that a worker might be trafficked, the combination of sophisticated anomaly detection and improved data science methods unlocks the value and knowledge in the data. In support of these goals, we critically analyze the risk analysis landscape as it pertains to the source of our data, advertisements associated with online voluntary commercial sex work and sex trafficking. In Chapter 2 we expand on the literature that intersects between the risk analysis and sex trafficking domains.

The crime of sex trafficking is complex and regularly misunderstood. Chapter 3's holistic evaluation of a major social disruption, the COVID-19 pandemic, on the U.S. sex trafficking system uniquely combines quantitative findings from anomaly detection with qualitative analysis of socioeconomic, political/governmental, and environmental factors that affect victims. We label this methodology as interdisciplinary data science because we apply domain expertise from law, public health, and public policy to guide decision-making through data analysis. Each step, from what we scraped to how we tested the data, was informed by the collective knowledge of the collaborators.

For example, in stratifying our data we incorporated expert insights from our informal interviews with antitrafficking professionals. Their insights fundamentally altered our methods and contradicted the underlying assumptions of previous attempts to measure sex trafficking through online commercial sex work advertisements. Our work identifies changes in the data that signify potential risk. We find that the dramatic decline of online commercial sex advertising associated with the pandemic meets with a confluence of factors that affects populations already in sex trafficking and those marginalized populations that may be pushed into sex trafficking. We also track the increasing trajectory of sex-trafficking-related advertising and hypothesize that it may continue to increase as social distancing restrictions lift across the country. To my knowledge, no other work currently exists that examines the impact of COVID-19 on independent commercial sex work or suspected sex trafficking.

This work is also salient to the present political and legislative climate. Although some research does exist using online sex advertisements as a proxy for prevalence, it was conducted before the April 2018 passage of the Fight Online Sex Trafficking Act—Stop Enabling Sex Trafficking Act (FOSTA-SESTA), the U.S. House and Senate legislation that effectively took down backpage.com (Allow, 2018) and other widely used online sex advertisement platforms. The ecosystem of sites on which sex traffickers and independent sex workers advertise has radically shifted; all of the data in this research analyzes advertisements posted after this important legislation, and thus offers a basis for comparison to previous work. This work makes a meaningful contribution to the literature with a first examination of online commercial sex advertising in a post–FOSTA-SESTA era.

Chapter 4 extends the time-series anomaly detection literature by crossing boundaries across multiple disciplines to demonstrate the utility of these techniques for the risk analysis community. We analyze findings from the application of machine learning, change point detection, and statistical process control on point, collective, and contextual anomalies and repeat the methods on a real-world suspected sex trafficking data set. This dissertation contributes to research not only by closing the quantitative literature gap in sex trafficking, but also by providing risk analysts with a starting point for their univariate time-series data. We establish which algorithms perform best for the different types of synthetic data, making our results helpful for risk analysts whose knowledge favors their domain expertise and data rather than specific methodologies.

5.2 Future Research Directions

Each chapter in the body of this dissertation has a distinct research focus, and each could be expanded upon. Directions of future research relevant to the discussion in Chapter 2 could include the development of a Bayesian Belief Network that incorporates the quantitative and qualitative insights of experts to calculate potential probabilities of entering sex trafficking. This type of graphical representation is especially appropriate for helping decision-makers visualize the impact of their policies. Similarly, Chapter 2 briefly discussed the merits of agent-based modeling in context of risk analysis. This would also be a potentially fruitful area of future research on sex trafficking. Modeling and simulating traffickers, victims, and law enforcement would allow for experimental testing of different types of policy levers.

An extension of Chapter 3’s research could include work on detecting the impact of COVID-19 (or any major social disruption) on other crimes that take advantage of vulnerable populations, such as domestic violence or hate crimes. Although data handling and sources will inevitably differ, this work would encourage the comprehensive analysis of other environmental, socioeconomic, and political factors that have an effect on social issues. Recent research details how crime rates have drastically fallen in many communities around the world, a change attributable to the government-mandated stay-at-home orders (Stickle & Felson, 2020). We know that many other crimes and social systems are affected by the pandemic, and future research can build on our work to measure that impact. Additionally, this research focuses specifically on the in-person subset of commercial sex work, but there exists a vast online presence of both independent commercial sex work and sex trafficking on websites such as Onlyfans.com. Media reports suggest that a large proportion of commercial sex has moved online due to the COVID-19 pandemic (Boseley, 2020). An extension of this work could measure the magnitude to which this presence has actually grown, critically analyzing methods to obtain this type of information, continuing to track the pandemic's impact, and detecting anomalies .

The anomaly detection algorithm comparison presented in Chapter 4 can be extended by incorporating more types of learning algorithms. Although we specifically experimented with the machine learning random forest methods, future work could include neural networks or statistical methods such as autoregressive integrated moving average (ARIMA) for the walk-forward validation. The augmentation of forecasting with machine learning methods for both change point detection and statistical process control were not covered in Chapter 4, but could also be a productive area of future research. Furthermore, this body of work specifically examines univariate time-series data, but future extensions could incorporate multivariate data. Along the same lines, this work could be enhanced by the aggregation with other data. Websites such as Rubmaps.com track online reviews for illicit massage businesses, and anomaly detection on more websites could also provide decision-makers with confidence that the effects of major social disruptions are being validated across multiple different platforms.

In conclusion, understanding the risks of sex trafficking provides an improved basis for decision-making about this murky and complex problem. This body of work helps unpack how activities surrounding sex trafficking introduce risk to communities. These contributions can have broader effects on how decision-makers, law enforcement, and antitrafficking professionals approach and understand the sex trafficking enterprise. Lastly, highlighting the lack of research in this domain and exploring data-driven methods reveals where more quantitative research needs to occur. In doing so, this body of work serves as a call to action for more scientific and evidence-based research into this devastating human problem.

Appendices

Appendix A

<u>Notes on Table 3.2</u>

Table 3.2 summarized statistics by city describing sex trafficking escort advertisement behavior over time. This data covered the period from January 2020 to May 2020, sorted in order of stay-at-home orders.

The following details how those statistics were calculated:

dip_velocity = $\frac{(\boldsymbol{lowest\ ad\ count - pre\ Covid\ avg})}{(\boldsymbol{lowest\ ad\ date - dip\ began})}$

lowest_ad_count = lowest raw number of advertisements

lowest_ad_date = date at which lowest_ad_count captured

increase_began = approximate date at which ad count began to rise again

depth(%) = $\frac{(\boldsymbol{lowest\ ad\ count})}{(\boldsymbol{pre\ Covid\ avg})} * \mathbf{100}$

Appendix B

Chapter 4 Data Description, Combined Flowchart for Algorithms, and More on Random Forest

The automated scrapers for this analysis retrieved 898,174 total advertisements in the United States over the period from July 11, 2019 to May 31, 2020. Because the scraper was running daily, however, many of these advertisements were duplications. If an advertiser were to place an ad to appear on Tuesday and then again on Friday, when the scraper would retrieve Friday's scrape the Tuesday advertisement would still appear, labeled for Tuesday. We prioritize the date the advertiser posted the advertisement and remove the duplicated advertisements to accurately characterize the actual daily advertising behavior, and to prevent inadvertently overweighting advertisements. This procedure reduced the data set significantly, which is represented in the data stratification of Appendix Figure 1.

These websites are littered with scammers attempting to gain credit card information (McCormick & Eberle, 2013). One co-author (Guikema) confirmed this by connecting with advertisers that listed only an email address and found they were all scammers. We removed advertisements that did not have phone numbers because "spam detection" research found that 78% of fake advertisements did not have phone numbers while 87% of real advertisements revealed phone numbers in the advertisement (Tran et al., 2011)

We stratify the data to separate likely sex trafficking advertisements from those likely for independent commercial sex workers. Although previous work uses movement or a "new to town" designation as indicators of trafficking, we believe that this labeling can be misleading because sex trafficking does not necessarily have to involve movement (Boecking et al., 2019; Ibanez &

Suthers, 2015). Instead we identify the clusters of phone numbers pointing to one single number and use that as a proxy for third-party-management. Our informal interviews with law enforcement, current and former sex worker, and anti-trafficking experts validated this notion and confirm this may be a more likely direct indicator. Although a generalized trafficking indicator, we acknowledge that it is not a perfect technique. We recognize that our method disregards the use of disposable phones that are common in nefarious practices, and we also do not include the myriad decoy advertisements law enforcement employs.

Appendix Figure 1 describes the stratification of the data as described in this section. The bottom layer includes an aggregated county by day of all of the previous layers information.

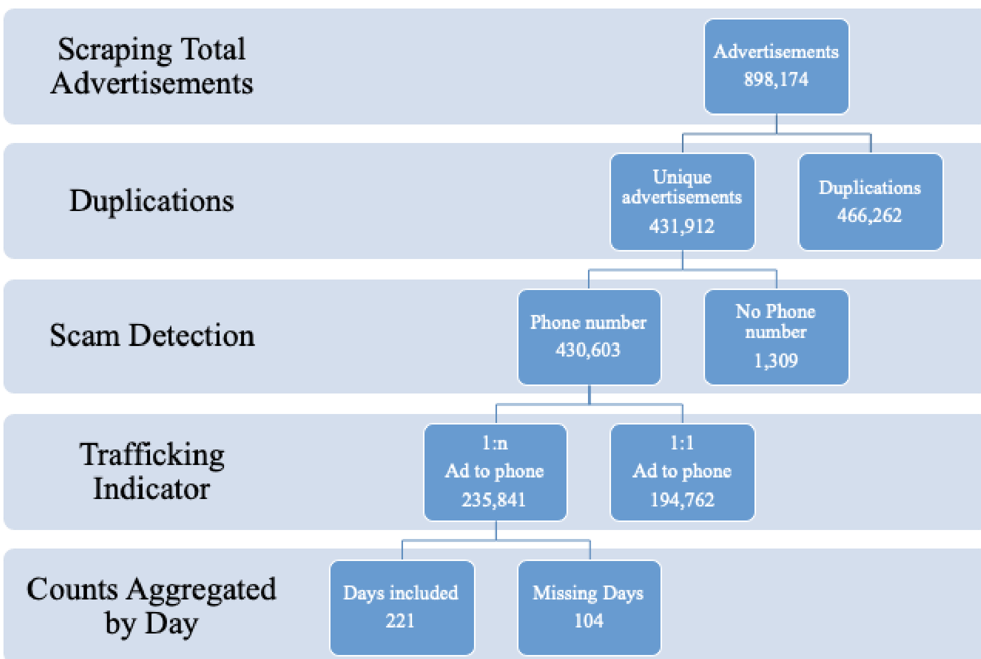

Appendix Figure 1: Data stratification for online advertising scraped off of rubratings.com from July 11, 2019 to May 31, 2020

Below in Appendix Table 1 are the descriptive statistics of data prior to the last step of aggregation. Although the far-left column describes the total number of advertisements considered in this analysis, the descriptive statistics of min, max, mean and standard deviation are all in reference to the counts of advertisements aggregated by day.

Appendix Table 1: Daily Descriptive Statistics for scraped advertisements (rubratings.com). July 2019 - May 2020

US Ad Count (total)	**Minimum (day)**	**Maximum (day)**	**Mean (day)**	**Std. Dev.(day)**
235,841	121	1671	987	363

Real-World Data Preprocessing Details

We used Appendix Table 2 data as a qualitative check for stationarity of the real-world sex trafficking data set. Our conclusion after acknowledging the large swings in mean in variance was to further pursue quantitative checks.

Appendix Table 2: Mean and variance calculated for daily count data split in half

Mean 1 = 1027.627273	Mean 2 = 948.162162
Variance 1 = 132068.633802	Variance 2 = 127828.928658

As mentioned previously, the data underwent an autocorrelation check using the lags across the time-series. We did this using the AutoCorrelation function native in the Statsmodel package for Python (Pertold et.al, 2019):

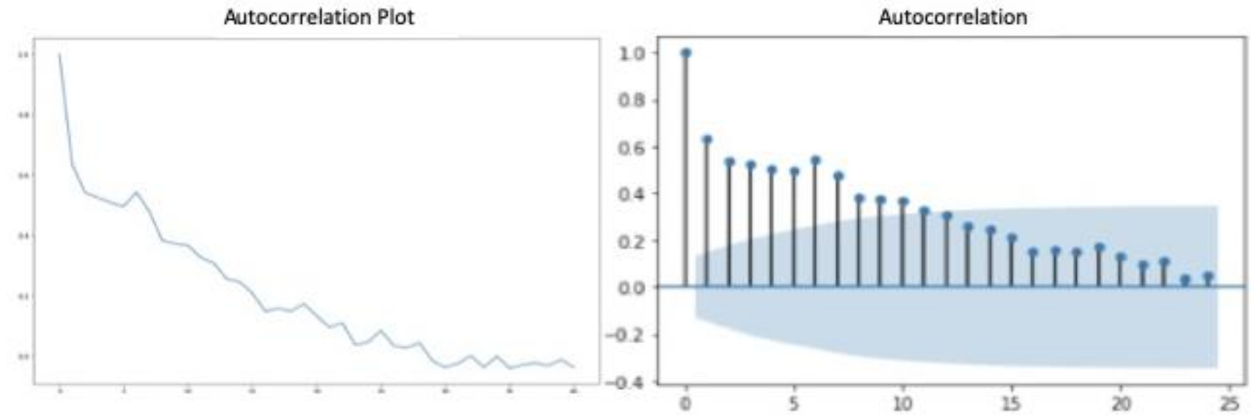

Appendix Figure 2: Autocorrelation plots. Left: autocorrelation of series Right: autocorrelation with lags

The autocorrelation helps determine if there is a correlation between the lagged values of the time-series. The observations establish a previous fixed point in time, and lags are also compared to that previous point in time. The confidence interval is shown by the blue cone, and the documentation of this package states that "correlation values outside of this cone are very likely a correlation and not a statistical fluke" (Pertold et al., 2019). Appendix Figure 2 shows a high correlation with the first lag and less correlation with the subsequent lags. The lags beyond 10 are within the cone and uncorrelated, and the series exhibits a fast decay meaning, that the future values do not have a very high correlation with its past values; this also suggests a conclusion of nonstationarity.

The final test for stationarity was conducted using the Augmented Dickey-Fuller test (Nielsen, 2016). In this case, the test failed to reject the null hypothesis that suggests that the time-series has a unit root and is nonstationary. Calculating a resultant p value of 0.4849, this test implies that the series has some time dependent structure.

In pursuit of appropriate data transformation, we used Unobserved Components Model (UCM) for seasonal decomposition, log transformation, and differencing. UCM can be summarized by the formula below (Harvey et.al, 1993; Perktold et.al, 2019):

$$y_t = \mu_t + \gamma_t + c_t + \sum_{j=1}^{k} \beta_{jx_{jt}} + \varepsilon_t$$

(trend). (seasonal) (cycle) (explanatory) (irregular)

The application of this statsmodel module makes it possible to include or exclude trends or seasonality, while also allowing for different types of levels that toggle between stochastic and deterministic components. It allows a researcher to tune in different levels, and in this case, we explored the use of a smoothing trend (integrated random walk), local linear trend, and local level (random walk; Perktold et.al, 2019).

The addition of each seasonality or trend component resulted in increased MAE (mean absolute error). This led to the conclusion that incorporating a seasonality or trend component would not improve the model performance. Below is a visualization of the resulting UCM output (Perktold et.al, 2019).

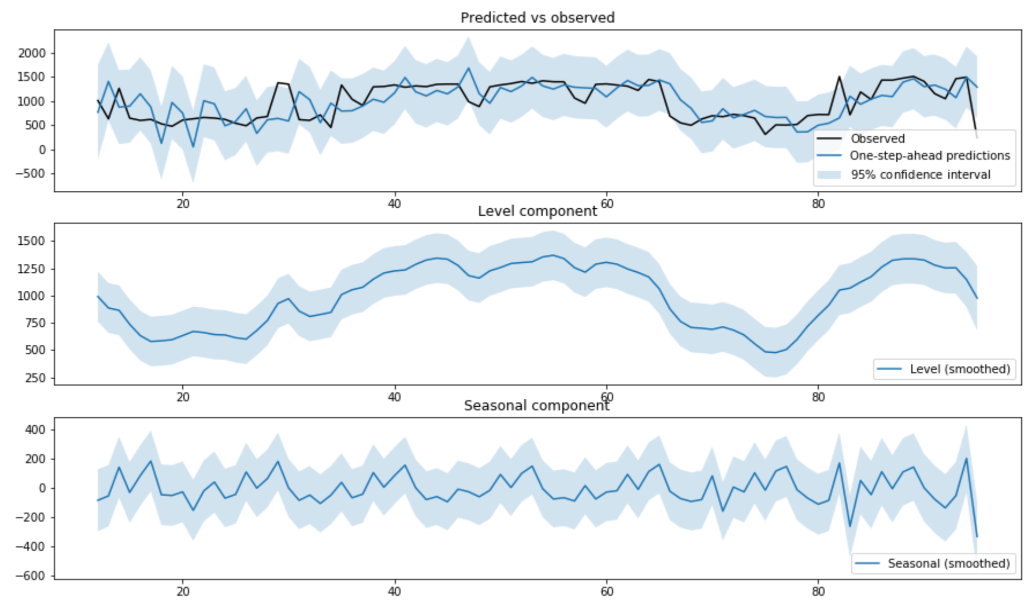

Note: The first 12 observations are not shown, due to approximate diffuse initialization.

Appendix Figure 3: Example of Predicted vs. Observed output from univariate Unobserved Components Model time-series analysis

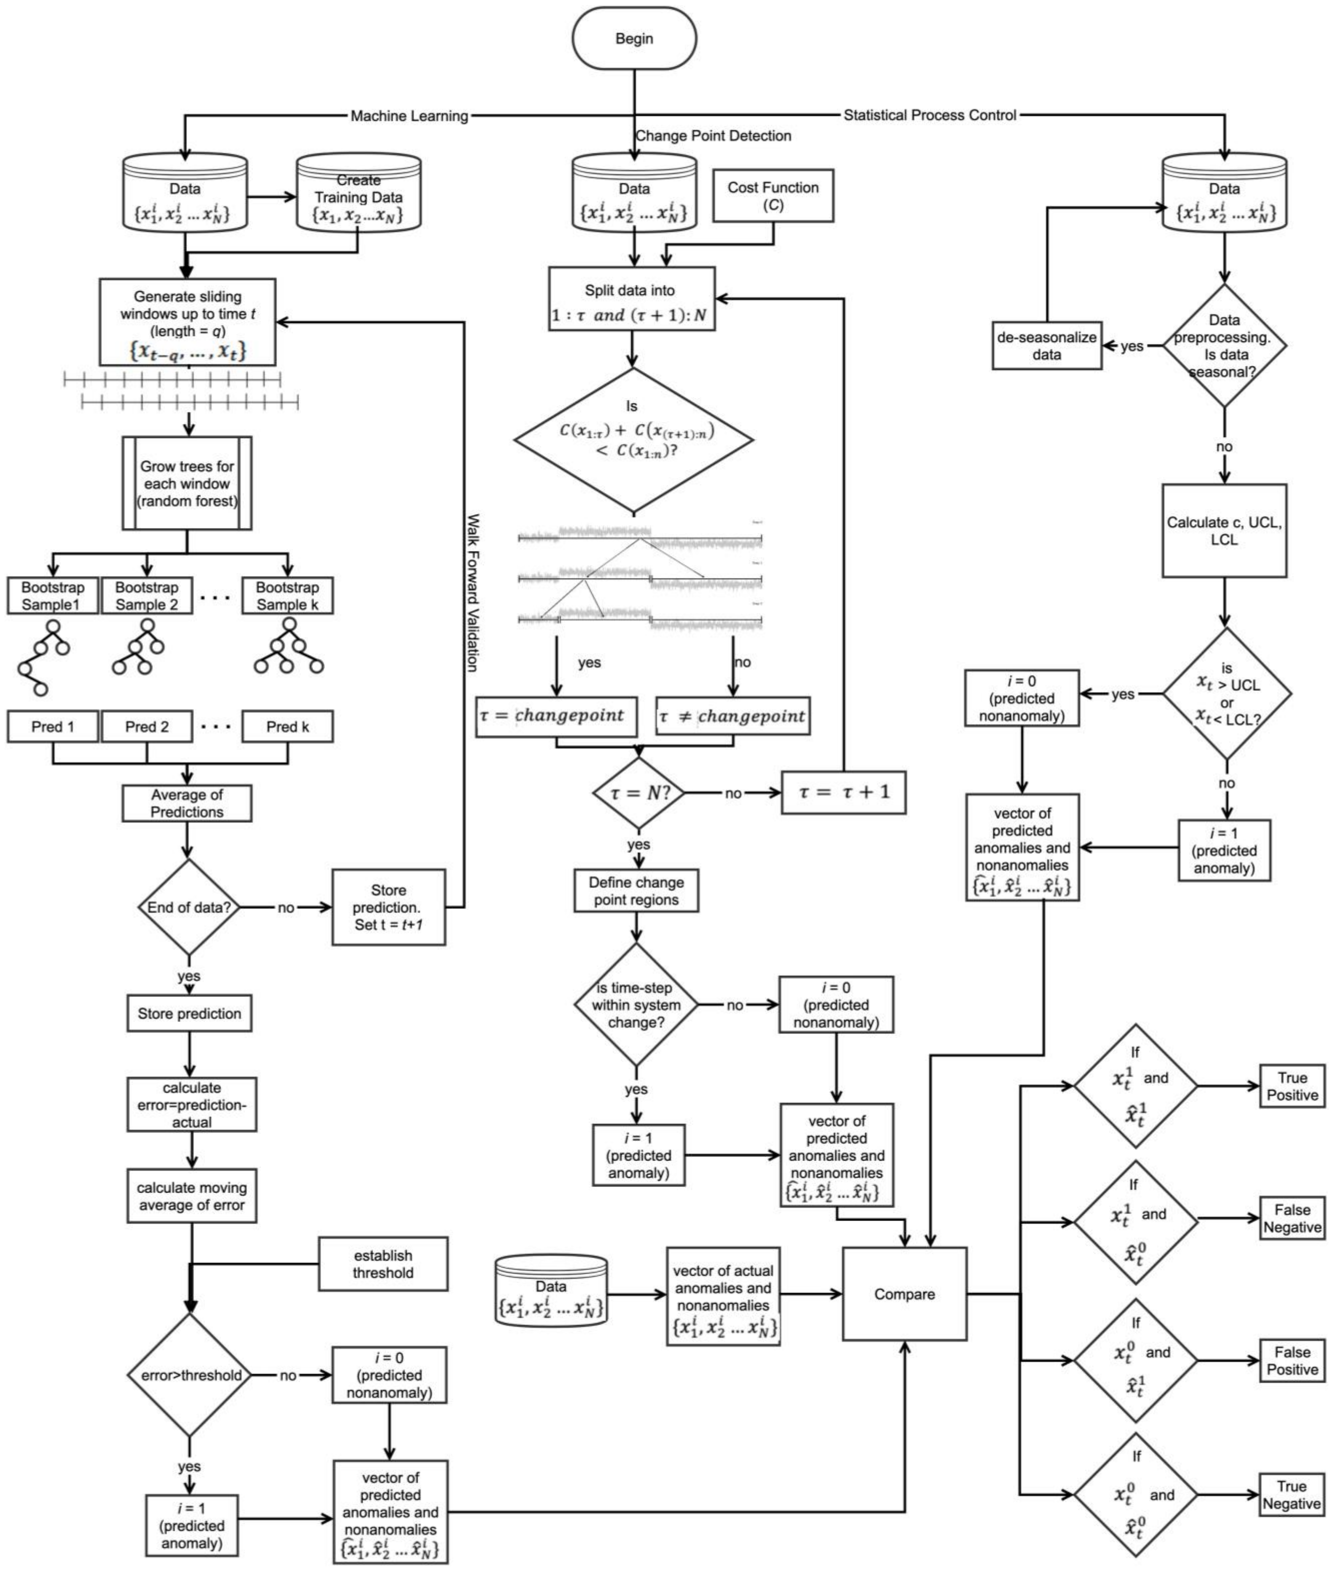

Appendix Figure 4: Combined flowchart of all algorithms and comparison methodology (Rodriguez-Galiano et al., 2016; Truong, 2020)

More on Random Forest

Recall from Section 4.3.2.1 among the many possible classical machine learning algorithms available for use in our walk forward validation, we use the random forest. Random forest is a form of decision tree, or algorithm in which entities or data points are categorized according to traits. A typical decision tree algorithm will generate rules that determine which features most strongly contribute to the final outcome. It is a modification of bagging or bootstrap aggregation (Friedman et al., 2001) using that same decision or regression tree many times on bootstrap-sampled versions of the training data (See Appendix Figure 4). Trees are generally high-variance, low-bias procedures that lend themselves well to bagging (Friedman et al., 2001). The algorithm randomly selects observations and only a portion of the feature space to grow an ensemble of trees, or a forest. The predictions from these trees are averaged making one model prediction.

One of the benefits of this technique is the averaging of the predictions which reduces the variance of notoriously noisy trees. As result of the random forest procedure, each tree is fit on slightly different data producing slightly different performance. Reducing the feature space (typically to one-third, Breiman 2001) at each split point in the tree to a random subset forces each tree in the forest to be different. The random forest uses unpruned trees making them slightly overfit to the training data which is desirable to produce less correlated predictions. It also has the desirable feature of avoiding overfitting predictions, in which a machine learning algorithm learns the training data too well and models random fluctuations in the data. If enough trees are used to generate the random forest, this will not happen. The random forest is one of the most popular and widely used machine learning techniques given the combination of its ease of use and its applications across a wide range of regression and classification problems.

Bibliography

Advisory Board. (2020). In these 22 states, the coronavirus epidemic is getting worse. https://www.advisory.com/daily-briefing/2020/06/15/coronavirus.

ACLU. (2020). *Is sex work decriminalization the answer?* https://www.aclu.org/sites/default/files/field_document/aclu_sex_work_decrim_research_brief.pdf

Ager, A. A., Kline, J. D., & Fischer, A. P. (2015). Coupling the biophysical and social dimensions of wildfire risk to improve wildfire mitigation planning. *Risk Analysis: An Official Publication of the Society for Risk Analysis*, *35*(8), 1393–1406.

Ahmed, A., & Seshu, M. (2012). "We have the right not to be 'rescued'": When anti-trafficking programmes undermine the health and well-being of sex workers. *Anti-Trafficking Review*, *1*, 149–168.

Ahmed, M., Mahmood, A. N., & Islam, M. R. (2016). A survey of anomaly detection techniques in financial domain. *Future Generations Computer Systems: FGCS*, *55*, 278–288.

Al-Rousan, S., Abuhussein, A., Alsubaei, F., Collen, L., & Shiva, S. (2020). *Ads-guard: Detecting scammers in online classified ads* [paper presentation]. IEEE Symposium Series on Computational Intelligence (SSCI), Canberra, Australia.

Albright, E., & D'Adamo, K. (2017). Decreasing human trafficking through sex work decriminalization. *AMA Journal of Ethics*, *19*(1), 122–26.

Allow States and Victims to Fight Online Sex Trafficking Act of 2017 to Communications Act of 1934, H.R. 1865, 115th Congress (2018). https://www.congress.gov/bill/115th-congress/house-bill/1865/text

Aminikhanghahi, S., & Cook, D. J. (2017). A survey of methods for time-series change point detection. *Knowledge and Information Systems*, *51*(2), 339–367.

Anthony, B. (2018). On-ramps, intersections, and exit routes: a roadmap for systems and industries to prevent and disrupt human trafficking. *Polaris Project*. polarisproject.org/a-roadmap-for-systems-and-industries-to-prevent-and-disrupt-human-trafficking.

Apostolakis, G. (1990). The concept of probability in safety assessments of technological systems. *Science, 250*(4986), 1359–64.

Aven, T. (2010). On how to define, understand and describe risk. *Reliability Engineering & System Safety*, *95*(6), 623–631.

Aupperlee, A. (2021, April 23). *Algorithm Uses Online Ads To Identify Human Traffickers*. Carnegie Mellon School of Computer Science. https://www.cs.cmu.edu/news/algorithm-uses-online-ads-identify-human-traffickers.

Bales, K. (2007). What predict human trafficking? *International Journal of Comparative and Applied Criminal Justice*, 31(2), 269-279.

Battaglia, J. (2021, March 27). Baltimore will no longer prosecute drug possession, prostitution and other low-level offenses. *CNN*. https://www.cnn.com/2021/03/27/us/baltimore-prosecute-prostitution-drug-possession/index.html

Benigni, M. C., Joseph, K., & Carley, K. M. (2017). Online extremism and the communities that sustain it: Detecting the ISIS supporting community on Twitter. *PloS One*, *12*(12), e0181405.

Benigni, M. C., Joseph, K., Carley, K. M. Online extremism and the communities that sustain it: Detecting the ISIS supporting community on Twitter. *PloS One*, *12*(12), e0181405.

Berger, M. (2020, April 28). Sex workers are falling through the cracks in coronavirus assistance programs around the world. *The Washington Post*. https://www.washingtonpost.com/world/2020/04/28/sex-workers-are-falling-through-cracks-coronavirus-assistance-programs-around-world/

Bier, V. M. (2007). Choosing what to protect. *Risk Analysis: An Official Publication of the Society for Risk Analysis*, *27*(3), 607–20.

Bigio, J., & Vogelstein, R. (2019). *Security implications of human trafficking*. Council on Foreign Relations. https://cdn.cfr.org/sites/default/files/report_pdf/Discussion_Paper_Bigio_Vogelstein_Security_Trafficking_OR.pdf

Bigio, J., & Vogelstein, R. (2019). *Women and terrorism*. Council on Foreign Relations. https://cfrd8-files.cfr.org/sites/default/files/report_pdf/Discussion_Paper_Bigio_Vogelstein_Terrorism_OR.pdf

Boecking, B., Miller, K., Kennedy, E., & Dubrawski, A. (2019). Quantifying the relationship between large public events and escort advertising behavior. *Journal of Human Trafficking*, *5*(3), 220–237.

Bontempi, G., Ben Taieb, S., & Le Borgne, Y.-A. (2013). Machine Learning Strategies for Time

Series Forecasting. In M.-A. Aufaure & E. Zimányi (Eds.), *Business Intelligence: Second European Summer School, eBISS 2012, Brussels, Belgium, July 15-21, 2012, Tutorial Lectures* (pp. 62–77). Springer Berlin Heidelberg.

Boseley, M. (2020, December 23). “Everyone and their mum is on it”: OnlyFans booms in popularity during the pandemic. *The Guardian*. https://www.theguardian.com/technology/2020/dec/23/everyone-and-their-mum-is-on-it-onlyfans-boomed-in-popularity-during-the-pandemic

Bouché, V., & Shady, S. (2017). A pimp’s game: A rational choice approach to understanding the decisions of sex traffickers. *Women & Criminal Justice*, 27(2), 91–108.

Braei, M., & Wagner, S. (2020). Anomaly detection in univariate time-series: a survey on the state-of-the-art. in *arXiv [cs.LG]* https://arxiv.org/pdf/2004.00433.pdf

Breiman, L. (2001). Random forests. *Machine Learning*, *45*(1), 5–32.

Breslin, S. (2020, 14 May). A sex worker reveals how the COVID-19 pandemic has changed sex work. *Forbes.com*. www.forbes.com/sites/susannahbreslin/2020/05/14/a-sex-worker-reveals-how-COVID-19-has-changed-sex-work/#1b49b0bb3a81.

Brownlee, J. (2020). Random forest for time series forecasting. *Machine Learning Mastery, 31*. machinelearningmastery.com/random-forest-for-time-series-forecasting/.

Caulkins, J. P., Kammer-Kerwick, M., Konrad, R., Maass, K. L., Martin, L., & Sharkey, T. (2019). A call to the engineering community to address human trafficking. *The Bridge on Cybersecurity, 49*(3). https://www.nae.edu/216534/A-Call-to-the-Engineering-Community-to-Address-Human-Trafficking

Chamberlain, L. (2018). FOSTA: A hostile law with a human cost. *Fordham Law Review*, *87*, 2171.

Chen, C., Dell, N., & Roesner, F. (2019). *Computer security and privacy in the interactions between victim service providers and human trafficking survivors* [paper presentation]. 28th ${USENIX} Security Symposium, Santa Clara, CA. https://www.usenix.org/conference/usenixsecurity19/presentation/chen

Cho, S.-Y., Dreher, A., & Neumayer, E. (2013). Does legalized prostitution increase human trafficking? *World Development*, *41*, 67–82.

Choi, K. R. (2015). Risk factors for domestic minor sex trafficking in the United States: A literature review. *Journal of Forensic Nursing*, *11*(2), 66–76.

Chu, S. K. H., & Glass, R. (2013). Sex work law reform in Canada: Considering problems with the Nordic model. *Alberta Law Review*, *51*, 101.

Clawson, H. (2006). Law enforcement response to human trafficking and the implications for victims: current practices and lessons learned. *US Department of Justice*. Retrieved from https://www.ncjrs.gov/App/abstractdb/AbstractDBDetails.aspx?id=238165.

Code of Federal Regulations: What Businesses Are Ineligible for SBA Business Loans?, 13 CFR § 120.110. (2017). https://www.law.cornell.edu/cfr/text/13/120.110

Cowin, R., Martin, H., & Stevens, C. (2020, July 27). Measuring evictions during the COVID-19 crisis. *Community Development Briefs*. https://www.clevelandfed.org/newsroom-and-events/publications/community-development-briefs/db-20200717-measuring-evictions-during-the-COVID-19-crisis.aspx.

Cox, L. A. Jr. (2009). Game theory and risk analysis. *Risk Analysis*, *29*(8), 1062–68.

Decrouez, G., & Robinson, A. (2013). Time-series models for border inspection data: Time-series models for inspection data. *Risk Analysis: An Official Publication of the Society for Risk Analysis*, *33*(12), 2142–2153.

De Gooijer, J. G., & Hyndman, R. J. (2006). 25 years of time series forecasting. *International Journal of Forecasting*, *22*(3), https://www.sciencedirect.com/science/article/pii/S0169207006000021

Deisler, P. F., Jr. (2002). A perspective: risk analysis as a tool for reducing the risks of terrorism. *Risk Analysis*, *22*(3), 405–413.

DeLateaur, M. (2016). From Craiglist to Backpage.com: Conspiracy as a strategy to prosecute third-party websites for sex trafficking. *Santa Clara Law Review*, *56*, 531–592.

Department of Justice (DOJ) (2012). *A National Overview of Prostitution and Sex Trafficking Demand Reduction Efforts, Final Report* (Document number: 238796). https://www.ncjrs.gov/pdffiles1/nij/grants/238796.pdf

Department of State. (2016). *Trafficking in persons report*. https://2009-2017.state.gov/documents/organization/258876.pdf

Dharmadhikari, A. S., Gupta, J., Decker, M. R., Raj, A., & Silverman, J. G. (2009). Tuberculosis and HIV: a global menace exacerbated via sex trafficking. *International Journal of Infectious Diseases: IJID: Official Publication of the International Society for Infectious Diseases*, *13*(5), 543–546.

Dubrawski, A. (2011). Detection of events in multiple streams of surveillance data. In C. Castillo-Chavez, H. Chen, W. B. Lober, M. Thurmond, & D. Zeng (Eds.), *Infectious disease informatics and biosurveillance: research, systems and case studies* (pp. 145–171). Springer US.

Ekberg, G. (2004). The Swedish law that prohibits the purchase of sexual services: best practices for prevention of prostitution and trafficking in human beings. *Violence Against Women*, *10*(10), 1187–1218.

Enders, C. K. (2010). *Applied missing data analysis*. Guilford Press.

Ezell, B. C., Bennett, S. P., von Winterfeldt, D., Sokolowski, J., & Collins, A. J. (2010). Probabilistic risk analysis and terrorism risk. *Risk Analysis, 30*(4), 575–89.

"The Facts." (2018, 9 November). *Polaris*. polarisproject.org/human-trafficking/facts.

Fahim, M., & Sillitti, A. (2019). Anomaly detection, analysis and prediction techniques in iot environment: a systematic literature review. *IEEE Access*, *7*, 81664–81681.

Farshchi, M., Weber, I., Della Corte, R., Pecchia, A., Cinque, M., Schneider, J., & Grundy, J. (2018). Contextual anomaly detection for a critical industrial system based on logs and metrics. *2018 14th European Dependable Computing Conference (EDCC)*, 140–143.

Feeney, H. (2013). RubMaps: Facilitating sexual exploitation in Chicago and beyond. CAASE. http://media.virbcdn.com/files/1b/3fb9ea6c4f55a345-RubmapsFinalReport.pdf

Flage, R., Aven, T., Zio, E., & Baraldi, P. (2014). Concerns, challenges, and directions of development for the issue of representing uncertainty in risk assessment. *Risk Analysis*, 34(7), 1196–207.

Fricker, R. D. (2013). *Introduction to statistical methods for biosurveillance: With an emphasis on syndromic surveillance*. Cambridge University Press.

Friedman, J., Hastie, T., Tibshirani, R., & Others. (2001). *The elements of statistical learning* (Vol. 1). Springer Series in Statistics New York.

Galizzi, M., & Tempesti, T. (2015). Workers' risk tolerance and occupational injuries. *Risk Analysis*, *35*(10), 1858–1875.

Gil, C., Rios Insua, D., & Rios, J. (2016). Adversarial risk analysis for urban security resource allocation. *Risk Analysis*, *36*(4), 727–741.

Goldstein, M., & Uchida, S. (2016). A comparative evaluation of unsupervised anomaly detection algorithms for multivariate data. *PloS One*, *11*(4), e0152173.

Gordon, E. (2017, August 3). Prostitution decriminalized: Rhode Island's experiment. *WHYY*. https://whyy.org/articles/prostitution-decriminalized-rhode-islands-experiment/

Gozdiak, Elzbieta M. (2011). *Data and research on human trafficking: Bibliography of research-based literature*. DIANE Publishing.

Goździak, E. M., Graveline, S., Skippings, W., & Song, M. (2015). Bibliography of research-based literature on human trafficking: 2008-2014. *Georgetown University.*

Greenberg, M., Cox, A., Bier, V., Lambert, J., Lowrie, K., North, W., Siegrist, M., & Wu, F. (2020). Risk Analysis: Celebrating the Accomplishments and Embracing Ongoing Challenges. *Risk Analysis: An Official Publication of the Society for Risk Analysis*, *40*(S1), 2113–2127.

Groff, E. R., Johnson, S. D., & Thornton, A. (2019). State of the Art in Agent-Based Modeling of Urban Crime: An Overview. *Journal of Quantitative Criminology*, *35*(1), 155–193.

Guikema, S. D. (2009). Game theory models of intelligent actors in reliability analysis: An overview of the state of the art. In V. M. M. Bier & M. N. Azaiez (Eds.), *Game Theoretic Risk Analysis of Security Threats* (pp. 13–31). Springer US.

Guikema, S. (2012). Modeling intelligent adversaries for terrorism risk assessment: some necessary conditions for adversary models. *Risk Analysis*, *32*(7), 1117–1121.

Harvey, A. C., & Jaeger, A. (1993). Detrending, stylized facts and the business cycle. *Journal of Applied Economics*, *8*(3), 231–247.

Hayes, M. A., & Capretz, M. A. M. (2014). Contextual anomaly detection in big sensor data. *2014 IEEE International Congress on Big Data*, 64–71.

Hill, D. J., & Minsker, B. S. (2010). Anomaly detection in streaming environmental sensor data:

A data-driven modeling approach. *Environmental Modelling & Software*. https://www.sciencedirect.com/science/article/pii/S1364815209002321?casa_token=5YsMm8qIqHwAAAAA:fQHRNBKAlgT7udmtnYiEc769L36EdYuKB0PK8QmzBo8nak0IQS8vTz0_m-rnGMkggkzxTYA

Howard, R. D. (2013). *The nexus of extremism and trafficking: Scourge of the world or so much hype?* https://www.socom.mil/JSOU/JSOUPublications/13-6_Howard_Nexus_FINAL.pdf

Huddleston, S. H., Learmonth, G. P., & Fox, J. (2008). Changing knives into spoons. *2008 IEEE Systems and Information Engineering Design Symposium*, 255–260.

Hughes, D. M. (2004). *Best Practices to Address the Demand Side of Sex Trafficking*. prostitutionetsociete.fr. http://www.prostitutionetsociete.fr/IMG/pdf/2004huguesbestpracticestoadressdemandside.pdf.

Hughes, D. M., et al. (2007). Modern-day comfort women: The US military, transnational crime, and the trafficking of women. *Violence Against Women*, *13*(9), 901–22.

Huisman, W., & Kleemans, E. R. (2014). The challenges of fighting sex trafficking in the legalized prostitution market of the Netherlands. *Crime, Law, and Social Change*, *61*(2), 215–228.

Ibanez, M., & Gazan, R. (2016). Virtual indicators of sex trafficking to identify potential victims in online advertisements. In J. Caverlee, R. Kumar, & H. Tong (Eds.), *Advances in Social Networks Analysis and Mining (ASONAM) 2016 IEEE/ACM International Conference* (pp. 818-824). IEEE.

Ibanez, M., & Suthers, D. D. (2014). *Detection of domestic human trafficking indicators and movement trends using content available on open internet source* [paper presentation]. 47th Hawaii International Conference on System Sciences (HICSS-47), Waikoloa, HI, USA. http://lilt.ics.hawaii.edu/papers/2014/Ibanez-Suthers-HICSS-2014.pdf

ILO (International Labour Organization). (2017). Global Estimates of Modern Slavery: Forced Labour and Forced Marriage. https://www.ilo.org/global/publications/books/WCMS_575479/lang--en/index.htm.

John Garrick, B., Hall, J. E., Kilger, M., McDonald, J. C., O'Toole, T., Probst, P. S., … Zebroski, E. L. (2004). Confronting the risks of terrorism: making the right decisions. *Reliability Engineering & System Safety*, *86*(2), 129–176.

Johnston, A., Friedman, B., & Sobel, M. (2015). Framing an emerging issue: how U.S. print and broadcast news media covered sex trafficking, 2008–2012. *Journal of Human Trafficking*, 1(3), 235–54.

Jones, L., Engstrom, D.W., Hilliard, T., & Diaz, M. (2007). Globalization and Human Trafficking. *Journal of Sociology and Social Welfare, 34*, 107-122.

Kaiser Family Foundation. (2020). State data and policy actions to address coronavirus. Retrieved June 3, 2020, from https://www.kff.org/health-costs/issue-brief/state-data-and-policy-actions-to-address-coronavirus/.

Kara, S. (2017). *Sex trafficking: Inside the business of modern slavery*. Columbia University Press.

Kasperson, R. E., Renn, O., Slovic, P., Brown, H. S., Emel, J., Goble, R., Kasperson, J. X., & Ratick, S. (1988). The Social Amplification of Risk: A Conceptual Framework. *Risk Analysis: An Official Publication of the Society for Risk Analysis*, *8*(2), 177–187.

Khamparia, A., Pande, S., Gupta, D., Khanna, A., & Sangaiah, A. K. (2020). Multi-level framework for anomaly detection in social networking. *Library Hi Tech*, *39*, 88.

Killick, R., Fearnhead, P., & Eckley, I. A. (2012). Optimal detection of changepoints with a linear computational cost. *Journal of the American Statistical Association*, *107*(500), 1590–1598.

Kim, Y. S., Rachev, S. T., Bianchi, M. L., Mitov, I., & Fabozzi, F. J. (2011). Time series analysis for financial market meltdowns. *Journal of Banking & Finance*, *35*(8), 1879–1891.

Kingston, S., & Thomas, T. (2019). No model in practice: a "Nordic model" to respond to prostitution?. *Crime, Law and Social Change,* 71, 423–439.

Kiss, L., Fotheringhame, D., & Mak, J. (2020). The use of Bayesian networks for realist evaluation of complex interventions: Evidence for prevention of human trafficking. *Of Computational Social Science*. https://link.springer.com/article/10.1007/s42001-020-00067-8.

Konrad, R. A., Trapp, A. C., Palmbach, T. M., & Blom, J. S. (2017). Overcoming human trafficking via operations research and analytics: Opportunities for methods, models, and applications. *European Journal of Operational Research*, *259*(2), 733–745.

Kroman, D. (2018, November 3). *Seattle's latest prostitution sting: Progressive or misguided?* Crosscut. https://crosscut.com/2018/11/seattles-latest-prostitution-sting-progressive-or-misguided

Lane, T., & Brodley, C. E. (1997). An application of machine learning to anomaly detection. *Proceedings of the 20th National Information Systems Security Conference*, *377*, 366–

380.

Laptev, N., Amizadeh, S., & Flint, I. (2015). Generic and scalable framework for automated time-series anomaly detection. *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1939–1947.

Larsen, Jacqueline, et al. *Delta 8.7*, Jan. 2019, delta87.org/wp-content/uploads/2019/01/Symposium_ModellingModernSlaveryRisk_Final.pdf.

Latonero, M. (2011). Human trafficking online: The role of social networking sites and online classifieds. *SSRN Electronic Journal*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2045851.

LCHT (Laboratory to Combat Human Trafficking). (2017, June 20). *Unpacking the myths: Human trafficking vs. human smuggling*. Laboratory to Combat Human Trafficking. https://combathumantrafficking.org/2017/06/unpacking-myths-human-trafficking-human-smuggling/

Lederer, L. J., & Wetzel, C. A. (2014). The health consequences of sex trafficking and their implications for identifying victims in healthcare facilities. *Annals of Health Law*, *23*(61), 61-91.

Levitin, G., & Hausken, K. (2010). Resource distribution in multiple attacks against a single target. *Risk Analysis*, *30*(8), 1231–39.

Li, S., Zhuang, J., & Shen, S. (2017). Dynamic forecasting conditional probability of bombing attacks based on time-series and intervention analysis: Conditional probability analysis of bombing attacks. *Risk Analysis*, *37*(7), 1287–1297.

Martin, L., & Hill, A. (2019). Debunking the myth of Super Bowl sex trafficking: Media hype or evidenced-based coverage? *Anti-Trafficking Review*, *13*, 13–29.

Matteson, D. S., & James, N. A. (2014). A nonparametric approach for multiple change point analysis of multivariate data. *Journal of the American Statistical Association*, *109*(505), 334–345.

McCarthy, B., Benoit, C., Jansson, M., & Kolar, K. (2012). Regulating Sex Work: Heterogeneity in Legal Strategies. *Annual Review of Law and Social Science*, *8*(1), 255–271.

McCormick, A. M., & Eberle, W. (2013). Discovering fraud in online classified ads [paper presentation]. 26th International FLAIRS Conference, St. Pete Beach, FL, USA. https://www.aaai.org/ocs/index.php/FLAIRS/FLAIRS13/paper/viewPaper/5928

McKinley, J. (2020). New York City region is now an epicenter of the coronavirus pandemic. *New York Times*. https://www.nytimes.com/2020/03/22/nyregion/Coronavirus-new-York-epicenter.html

Merrick, J., & Parnell, G. S. (2011). A comparative analysis of PRA and intelligent adversary methods for counterterrorism risk management. *Risk Analysis*, *31*(9), 1488–1510.

Miller, R. G., Jr. (2012). *Simultaneous statistical inference.* Springer Science & Business Media.

Montgomery, D. C. (2012). *Statistical quality control*. Wiley Global Education.

Mushtaq, R. (2011). *Augmented Dickey Fuller Test*. https://doi.org/10.2139/ssrn.1911068

Nagpal, C., Miller, K., Boecking, B., & Dubrawski, A. (2015). An entity resolution approach to isolate instances of Human Trafficking online. *arXiv*. http://arxiv.org/abs/1509.06659.

Nielsen, A. (2016). *Time series analysis with Python*. https://github.com/AileenNielsen/TimeSeriesAnalysisWithPython

Okeowo, A. (2020). The fragile existence of sex workers during the pandemic. *The New Yorker*. https://www.newyorker.com/news/news-desk/the-fragile-existence-of-sex-workers-during-the-pandemic

Omar, S., Ngadi, A., & Jebur, H. H. (2013). Machine learning techniques for anomaly detection: an overview. *International Journal of Computer Applications in Technology*, *79*(2). http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.402.6779&rep=rep1&type=pdf

Palaniappan, A., Bhargavi, R., & Vaidehi, V. (2012). Abnormal human activity recognition using SVM based approach. *2012 International Conference on Recent Trends in Information Technology*, 97–102.

Parpoula, C., Karagrigoriou, A., & Lambrou, A. (2017). Epidemic intelligence statistical modelling for biosurveillance. *Mathematical Aspects of Computer and Information Sciences*, 349–363.

Paté-Cornell, E., & Guikema, S. (2002). Probabilistic modeling of terrorist threats: A systems analysis approach to setting priorities among countermeasures. *Military Operations Research*, 7(4), 5–23.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., & Dubourg, V. (2011). Scikit-learn: Machine learning in Python. *The Journal of Machine Learning Research*, *12*, 2825–2830.

Polaris Project. (2020, September 24). *Separating fact from fiction: Recent cases of recovered missing children and what they show about child sex trafficking*. Polaris. https://polarisproject.org/blog/2020/09/separating-fact-from-fiction-recent-cases-of-recovered-missing-children-and-what-they-show-about-child-sex-trafficking/

Polaris Project. (2020, November 12). *2019 U.S. National Human Trafficking Hotline Statistics*. https://polarisproject.org/2019-us-national-human-trafficking-hotline-statistics/

Raymond, J. G. (2004). Ten reasons for not legalizing prostitution and a legal response to the demand for prostitution. *Journal of Trauma Practice*, *2*(3-4), 315–32.

Riegler, A. (2007). Missing the mark: Why the trafficking victims protection act fails to protect sex trafficking victims in the United States. *Harvard Journal of Law & Gender*, *30*, 231.

Rios, J., & Insua, D. R. (2012). Adversarial risk analysis for counterterrorism modeling. *Risk Analysis*, *32*(5), 894–915.

Rohrbeck, C. (2013). *Detection of changes in variance using binary segmentation and optimal partitioning*. lancaster.ac.uk. https://www.lancaster.ac.uk/pg/rohrbeck/ResearchTopicI.pdf

Rodriguez-Galiano, V. F., Sanchez-Castillo, M., Dash, J., Atkinson, P. M., & Ojeda-Zujar, J. (2016). Modelling interannual variation in the spring and autumn land surface phenology of the European forest. Biogeosciences , 13(11), 3305–3317.

Romilly, P. (2007). Business and climate change risk: a regional time series analysis. *Journal of International Business Studies*, *38*(3), 474–480.

Samuel, A., & Guikema, S. D. (2012). Resource allocation for homeland defense: Dealing with the team effect. *Decision Analysis*, *9*(3), 238–52.

Sandler, T., & Arce, D. G. (2003). Terrorism & game theory. *Simulation & Gaming*, *34*(3), 319–337.

Savage, D., Zhang, X., Yu, X., Chou, P., Wang, Q. (2014). Anomaly detection in online social networks. *Social Networks*, *39*, 62–70.

Savona, E. U., and S. Stefanizzi. (Eds.). (2007). *Measuring human trafficking: Complexities and pitfalls*. Springer.

Schwarz, C., Xing, C., Daugherty, R., Watt, S., & Britton, H. E. (2020). Frontline workers' perceptions of human trafficking: Warning signs and risks in the Midwest. *Journal of Human Trafficking*, *6*(1), 61–78.

Scott, A. J., & Knott, M. (1974). A cluster analysis method for grouping means in the analysis of variance. *Biometrics*, *30*(3), 507–512.

Seabold, S., & Perktold, J. (2010). Statsmodels: Econometric and statistical modeling with python. *Proceedings of the 9th Python in Science Conference*, *57*, 61.

Shao, Q., Liu, H., Zhu, P., Wei, W., Zhou, H., & Yang, M. (2020). The effects of key influencing factors on unsafe events in airport flight areas according to time series data. *Risk Analysis*. https://doi.org/10.1111/risa.13609

Shelley, L. (2012). The relationship of drug and human trafficking: A global perspective. *European Journal on Criminal Policy and Research*, *18*(3), 241–253.

Shewart, W. A. (1931). Economic Control of Quality of Manufactured Product ASQC (Milwaukee). *Revised Statutes: Under Arrangement of the Official Louisiana Revised Statutes of 1950. Louisiana*, 1980.

Shively, M., Kliorys, K., Wheeler, K., & Hunt, D. (2012). A national overview of prostitution and sex trafficking demand reduction efforts, final report. *National Institute of Justice.* https://www.ncjrs.gov/pdffiles1/nij/grants/238796.pdf

Smith, A., & Cockayne, J. (2020, April 2). The impact of COVID-19 on modern slavery. *Our World.* https://delta87.org/2020/03/impact-COVID-19-modern-slavery

Stickle, B., & Felson, M. (2020). Crime rates in a pandemic: The largest criminological experiment in history. *American Journal of Criminal Justice,* 1–12.

Szekely, P., Knoblock, C. A., Slepicka, J., Philpot, A., Singh, A., Yin, C., Kapoor, D., Natarajan, P., Marcu, D., Knight, K., Stallard, D., Karunamoorthy, S. S., Bojanapalli, R., Minton, S., Amanatullah, B., Hughes, T., Tamayo, M., Flynt, D., Artiss, R. … Ferreira, L. (2015). *Building and using a knowledge graph to combat human trafficking* [paper presentation].

Proceedings of the 14th International Semantic Web Conference (ISWC 2015), Bethlehem, PA, USA. https://usc-isi-i2.github.io/papers/szekely15-iswc.pdf

Talagala, P. D., Hyndman, R. J., Smith-Miles, K., Kandanaarachchi, S., & Muñoz, M. A. (2020). Anomaly detection in streaming nonstationary temporal data. *Journal of Computational and Graphical Statistics*, *29*(1), 13–27.

Taylor, S. J., & Letham, B. (2018). Forecasting at scale. *The American Statistician*, *72*(1), 37–45.

Todres, J., & Diaz, A. (2020, September 21). COVID-19 and human trafficking—the amplified impact on vulnerable populations. *JAMA Pediatrics*. https://jamanetwork.com/journals/jamapediatrics/fullarticle/2770536

Trafficking Victims Protection Act (TVPA) of 2000, Pub. L. No. 106-386, 114 Stat. 1466 (codified as amended in scattered sections of 8, 18, 22 U.S.C.).

Tran, H., Hornbeck, T., Ha-Thuc, V., & Cremer, J. (2011). Spam detection in online classified advertisements. *Joint WICOW/AIRWeb.* https://dl.acm.org/doi/abs/10.1145/1964114.1964122?casa_token=fYdCy7j_T2MAAAAA:tns9jAcl5vzmPg5xUfArLPWyWf_KfoFGf4q3thexqfBPuL7_8L2lKd0NGbDWTVTPCyFY-x-MiRU).

Tran, H., Hornbeck, T., Ha-Thuc, V., Cremer, J., & Srinivasan, P. (2011). Spam detection in online classified advertisements. In Proceedings of the 2011 Joint WICOW/AIRWeb Workshop on Web Quality (WebQuality '11). ACM, New York, NY, USA, 35–41.

Truong, C., Oudre, L., & Vayatis, N. (2020). Selective review of offline change point detection methods. *Signal Processing*, *167*, 107299.

Tyldum, G., & Brunovskis, A. (2005). Describing the unobserved: Methodological challenges in empirical studies on human trafficking. *International Migration* , *43*(1-2), 17–34.

United Nations Security Council. (n.d.). *Identifying and exploring the nexus between human trafficking, terrorism, and terrorism financing*. Counter-Terrorism Committee Executive Directorate.

United States Department of Justice, Civil Rights Division. (2006). Report on activities to combat human trafficking. *Report on Activities to Combat Human Trafficking*.

United States of America: Victims of Trafficking and Violence Protection Act of 2000, Public Law 106-386 [H.R.3244]. (2000, October 28). https://www.refworld.org/docid/3ae5b6104.html

van den Burg, G. J. J., & Williams, C. K. I. (2020). An evaluation of change point detection algorithms. in *arXiv [stat.ML]*. arXiv. http://arxiv.org/abs/2003.06222

Vuolajärvi, N. (2019). Governing in the name of caring—the Nordic model of prostitution and its punitive consequences for migrants who sell sex. *Sexuality Research and Social Policy, 16*, 151–165.

Weitzer, R. (2015). Researching prostitution and sex trafficking comparatively. *Sexuality Research & Social Policy*, 2015.

Welsh-Huggins, A. (2017, October 28). Buy sex, go to school: Anti-trafficking classes target johns. *The Seattle Times*. www.seattletimes.com/nation-world/buy-sex-go-to-school-anti-trafficking-classes-target-johns/.

Wilson, R. S., Winter, P. L., Maguire, L. A., & Ascher, T. (2011). Managing wildfire events: risk-based decision making among a group of federal fire managers: Managing wildfire events. *Risk Analysis*, *31*(5), 805–818.

Worsnop, C. Z. (2019). The disease outbreak-human trafficking connection: A missed opportunity. *Health Security*, *17*, 181–192.

Wu, H. (2016). A survey of research on anomaly detection for time series. *2016 13th International Computer Conference on Wavelet Active Media Technology and Information Processing (ICCWAMTIP)*, 426–431.

Zimmerman, C., Hossain, M., Yun, K., Gajdadziev, V., Guzun, N., Tchomarova, M., … Watts, C. (2008). The health of trafficked women: a survey of women entering posttrafficking services in Europe. *American Journal of Public Health*, *98*(1), 55–59.

Zimmerman, C., Hossain, M., & Watts, C. (2011). Human trafficking and health: a conceptual model to inform policy, intervention and research. *Social Science & Medicine*, *73*(2), 327–335.