

CRIME AND JUSTICE BULLETIN

NUMBER 260 | SEPTEMBER 2023

Text mining police narratives of domestic violence events to identify coercive control behaviours

Min-Taec Kim and George Karystianis^a^a School of Population Health, University of New South Wales

AIM

To construct a measure of coercive control behaviours from free-text narratives recorded by the NSW Police Force for each domestic violence event, and to assess how useful this measure might be in detecting coercive control behaviours when combined with existing measures and/or predicting which events are followed by violence within 12 months.

METHOD

We developed a text-mining system to capture a (non-exhaustive) set of behaviours that could be used to exert coercive control. Our text mining system consisted of a set of rules and dictionaries, with our definition of 'coercive control behaviours' drawn from Stark (2007). We concentrated on behaviours that were not already well captured by fixed fields in the police system. We applied this text mining system across all police narratives of domestic violence events recorded between 1 January 2009 and 31 March 2020, and used this data to construct our measure of coercive control behaviour. We then compared the incidence of coercive control behaviours using our measure against existing incident categories (where they exist). Finally, we estimated a gradient boosted decision tree (xgboost) model three times (once with standard predictors, once with just our coercive control measure, and once with both sets of variables included) and compared the performance of each model. This allowed us to estimate the contribution of our measure for improving the prediction of future violence.

RESULTS

There were 852,162 behaviours extracted by the text-mining system across 526,787 domestic violence events, with 57% of these events having at least one coercive control behaviour detected and 8% having three or more distinct subcategories of coercive control behaviours. These events were associated with 223,645 unique persons of interest. Our text mining system agreed with the pre-existing police incident categories (where they exist) in 80 to 99% of cases, with the text mining approach identifying an additional 30 to 60% of events that were not identified using the police incident categories, depending on the behaviour. The inclusion of the text mining variables did not improve the performance of our predictive model.

CONCLUSION

Our text mining system successfully extracted coercive control behaviours from police narratives, allowing us to identify how often these events are recorded in police narratives and supplement existing police categories of DV behaviours. Use of our measure in conjunction with existing incident categories significantly expands our ability to identify coercive control behaviours, however it does not improve our ability to predict which events are followed by further domestic violence.

KEYWORDS

Domestic violence

intimate partner violence

coercive control

abuse

text mining

INTRODUCTION

On 21 October 2020, the Joint Select Committee on Coercive Control was appointed by the Parliament of New South Wales (NSW) to inquire and report on coercive control in domestic relationships. The first recommendation from this Committee was the development of a stand-alone offence for coercive control, a recommendation that was supported by the NSW Government and passed into legislation by NSW Parliament on 16 November 2022.

The introduction of a stand-alone offence was justified based on both the serious harm arising from coercive control behaviours themselves, as well as the possibility that charges of this offence could prevent serious domestic violence in the future. One of the key pieces of evidence considered by the Committee (Parliament of New South Wales, 2021) was that in 111 of the 112 intimate partner homicide cases reviewed by the NSW Domestic Violence Death Review Team (DVDRT) the relationship between the abuser and the victim was characterised by coercive control (Domestic Violence Death Review Team, 2020). It was argued to the joint select committee by that the existence of coercive control legislation may have prevented some of these deaths by deterring or incapacitating the abuser before the violence escalated.¹

While the high prevalence of coercive control in intimate partner homicides is well recognised, much less is known about the prevalence of coercive control in non-fatal domestic violence incidents and the extent to which coercive control can be used operationally as a predictor of future violent offending. A key reason for this is the difficulty observing, identifying, and constructing a measure of coercive control that can be practically captured and operationalised by police. The methods used to coerce and control a victim can be subtle, deeply context specific, purposefully chosen to avoid detection, and may not constitute crimes in and of themselves. Isolated coercive control behaviours may not be recognised by family, friends or even the victim as an issue. This makes identifying coercive control behaviours challenging for police and creates fundamental problems for data collection and analysis.

Defining coercive control

Coercive control in domestic violence describes an abusive pattern or course of conduct, predominantly perpetrated by men against women, which aims to exert domination and control over the other party in an intimate partner relationship. The term coercive control was first coined by Susan Schechter in the early 1980s as part of her work in the *battered women's movement* (Kuennen, 2007) with related ideas being developed and disseminated by The Duluth Domestic Abuse Intervention Project (most famously through the 'Power and Control Wheel') as part of their broader community-based response. Coercive control is closely related to *intimate terrorism* (Johnson, 2008), *abusive-controlling violence* (Jaffe et al., 2008) and other terms in the domestic violence literature that describe gendered patterns of domination.

Although there are numerous definitions of coercive control in the literature, three features are common across most of these definitions (Hamberger et al., 2017):

1. Intentionality and motivation within the abuser to obtain control over the victim;
2. Perception of the behaviour as negative by the victim;
3. The ability of the perpetrator to make a credible threat.

As coercive control is defined in terms of both the intention of the abuser and the impact on the victim, this pattern can be realised using a range of abusive behaviours across a wide range of domains (e.g. physical, psychological, emotional) and is likely heavily dependent on the specific context of each relationship. Stark (2007) identifies violence, intimidation, isolation, and control as the defining features

¹ For police to prevent domestic violence homicides through the use of a coercive control offence, they need to be aware of events prior to the domestic violence incident – something that is not true of most lethal or near-lethal domestic violence events (Sherman, 2018). If the introduction of this offence does not expand the proportion of homicides with prior contact with the police, then the introduction of the offence will have limited scope to reduce the incidence of homicides.

of coercive control, and provides a set of behaviours that may be utilised in coercive control as seen in Table 2. This set of behaviours form the basis of the rules and dictionaries that constitute the text mining system in the current study.

Previous research

There are two major strands in the literature on intimate partner violence (Johnson, 1995; Myhill, 2015): the *family violence perspective* and; the *feminist perspective*. Studies from the family violence perspective (e.g. Straus, 1979; Straus, 2011) emerged from analyses of national survey data and often emphasise physical abuse and the gender symmetry in self-identified perpetration of physical violence, while acknowledging the large difference in harm experienced by women. Conversely, studies from the feminist perspective (e.g. Dobash & Dobash, 1979; Stark, 2007) started from a focus on trying to understand the experience of women in violent relationships, emphasising the existence of traditional gender roles and patriarchal power structures, and the role of coercive and controlling behaviours in reinforcing these ideologies within the relationship. As Wiener (2017) puts it, “the heart of the distinction between the feminist school’s and the family violence researchers’ descriptions of intimate partner abuse is the presence or absence of what Evan Stark has labelled ‘coercive control’” (p. 500).

These competing viewpoints have led to the development of a dual typology of domestic violence (Johnson 1995; Johnson 2008). One type (*situational couple violence*) is less serious, gender symmetric and does not involve coercive control, and the other, potentially more serious type (‘intimate terrorism’), is predominantly perpetrated by men towards women and is characterised by coercive control.² Although there is widespread acceptance of this dual typology for heterosexual couples, there is no such consensus in the literature for those categorising domestic violence in lesbian, gay, bisexual, transgender and queer relationships (Donovan & Hester, 2014; Hassouneh & Glass, 2008), with only a small literature applying these categories to same-sex relationships (Frankland and Brown, 2014).

Measuring the prevalence of coercive control

The most common approach for identifying coercive control is to measure the self-reported prevalence of specific behaviours related to coercive control that were either experienced by a victim and/or committed by a perpetrator. Researchers then define a cut-off value for the number of behaviours above which an individual has committed coercive control. Other approaches use a k-means clustering method to choose a threshold value that separates individuals into a low control or high control group, using measures of experience of control directly, or the extent to which coercion defines interactions within the relationship (Johnson & Leone, 2005; Hardesty et al., 2015).

Hamberger et al. (2017) find 22 different definitions of coercive control in the broader literature, and find that existing measures rarely capture all components of the definition of coercive control as these components are difficult to observe and quantify (i.e. the intentions of the abuser and/or the credibility of their threats). As definitions of coercive control are conceptually complex and require context to identify target behaviours, measures often present a trade-off between feasibility and accuracy.³ This suggests that different measures may be required in different contexts to identify prevalence and categorise ongoing patterns of domestic abuse.

The wide range of definitions, measures and samples considered in the literature has resulted in a correspondingly wide range of estimates of the prevalence of coercive control, making comparisons between studies and extrapolation to other settings challenging. Michalski (2005) concludes that 43% of violent male partners are intimate terrorists from data extracted from the General Social Survey in Canada. This is significantly higher than estimated by Johnson and Leone (2005) who report that 35% (81/230) of violent male partners are intimate terrorists based on telephone survey data in the United States of America (U.S.).

² Walby and Towers (2018) reconcile these perspectives by arguing that escalation is a function of lacking resources, arguing that domestic violent crime can be seen as a subset of violent crime that has gendered dimensions, rather than having two stable forms.

³ Clinical interviews with both partners represent the most nuanced, least feasible option, while relying on existing data collection being the most feasible but least nuanced. Structured interviews and self-reported scales in surveys may be significantly more feasible, but subject to methodological constraints that make them unsuitable for specific research or operational requirements.

Measures that include partners who use coercive control behaviours through exclusively non-violent behaviours may detect even higher rates (Walby & Myhill, 2001), while much lower rates are found when measures are restricted to severe harm and/or violent actions. For example, in a review of evidence from eight studies, Straus (2011) finds that 17% of violent male partners are intimate terrorists. In an Australian study, Boxall and Morgan (2021) present experiences of coercive control from an online survey of 15,000 women aged 18 years and older. They report that 11% of respondents had experienced three or more coercive control behaviours⁴ in the three months prior to the survey, with jealousy, monitoring of movements, financial abuse, social isolation, emotional abuse and/or threats being the most common behaviours.

Very few studies have considered the feasibility of a measure of coercive control that is not based on self-report. One exception is Barlow (2020), who considered the extent to which body-worn cameras could be used to capture coercive control behaviours, in the context of securing convictions for a coercive control offence in the United Kingdom. She finds that the ability of body-worn cameras to collect evidence is significantly hindered by the discretionary use of cameras, with police taking footage only at more acute incidents, and the (often) hidden nature of coercive control behaviours.⁵ These issues echo many of the concerns with data sources that are based on police collection such as the data source in the current study.

Coercive control as a predictor of significant harm

One key motivation for understanding when and where coercive control is taking place is the potential for it to be used as a predictor for homicide and other severe violent offences in domestic relationships. Several studies have found that coercive controlling behaviours are useful as predictors of further non-fatal, domestic violence offending, but the evidence on the prediction of more severe offending is limited.

Myhill and Hohl (2019) argue that coercive control is the *golden thread* running through risk identification and that prediction models will need to allow for clusters of risk factors (rather than estimating the predictiveness of specific behaviour) to be effective. Aguilar Ruiz and González-Calderón (2020) find that death threats and degrading treatment increase the risk of victimisation, while in the NSW context, Rahman (2018) finds that emotional abuse is a key predictor of revictimisation in her study of 336 individuals who reported experiencing intimate partner violence.

Conversely, an analysis of the NSW Domestic Violence Safety Assessment Tool (Leung & Trimboli, 2022), using 234,454 incidents between intimate partners, found that the measure of coercive control included in the tool contributed little to the prediction of repeat victimisation. This is consistent with a larger gap in the literature on the effectiveness and feasibility of Intimate Partner Violence (IPV) specific risk measures (Graham et al., 2021), whether risk assessment instruments are more effective at predicting different types of offending (e.g. stalking versus more severe violence incidents such as investigated in Gerbrandij et al., 2018) and the accuracy of practitioners' IPV risk assessments in general (Svalin & Levander, 2020).

Legislation introduced across jurisdictions

The difficulties defining and measuring coercive control have not deterred multiple jurisdictions from creating new criminal offences that target coercive control behaviours. The criminalisation of coercive control has been, at least in part, driven by the perception that existing offences are too narrowly defined to prosecute many abusers (Stark, 2018). The introduction of a new offence creates significant issues for legislators - if the scope is too narrow the offence may be infrequently used, but if the scope is too broad

4 The 13 behaviours here were threatening or abusing them online or using technology (e.g. over the phone or on social media); stalking them online or in person; constantly insulting them to make them feel ashamed, belittled or humiliated, or shouting, yelling or verbally abusing them to intimidate them; damaging, destroying or stealing their property; threatening to hurt their family, friends, children and/or pets; the perpetrator threatening to hurt themselves; monitoring their time and making them account for their whereabouts; using their money or shared money or making important financial decisions without talking to them; being jealous or suspicious of their friends; accusing them of having an affair; interfering with their relationship with other family members; preventing them from doing things to help themselves (e.g. going to medical appointments, taking medication); and restricting their use of their phone, the internet or the family car.

5 The role of body worn cameras in domestic violence cases is highly contested in the literature, including the impact on the outcome of cases and the well-being of victims (Harris, 2020). See Pfitzner, Walklate and McCulloch (2022) for further discussion of the difficulties of implementing body-worn cameras in domestic violence.

there may be uncertainty about which behaviours will be enforced under the new regime. This uncertainty then makes it difficult to estimate the impact of criminalisation, as this will depend on which events are ultimately charged under the new offence (e.g. how the *course of conduct* or *pattern of behaviour* is defined and which behaviours are deemed to be abusive (Barlow et al., 2018; Brennan, Burton, Gormally & O’Leary, 2018; Walklate et al., 2018)).

One of the explicit aims of expanding the scope of behaviours by criminalising coercive control is to prompt and/or enable police to intervene in domestic settings where previously they would not have intervened (Barlow et al., 2018; Whalley & Hackett, 2017). There is an implicit assumption that increasing policing will lead to reduced rates of domestic violence. However, this assumption has been criticised by many scholars, who argue that increased policing provides limited benefits (e.g., a sense of safety through the incarceration and/or deterrence of offenders) to only a small group of women⁶ while negatively impacting vulnerable groups. Queer people, people with disabilities and Aboriginal and Torres Strait islander peoples may be negatively impacted due to differences in how these groups are policed, how offences are enforced, and how these communities are considered by the criminal justice system (Walklate & Fitz-Gibbon, 2021; Watego et al., 2021; Whalley & Hackett, 2017).

A summary of specific reforms and offences that have been introduced over the last 10-15 years are described in Table 1 below.⁷

Table 1. Summary of coercive control legislation across jurisdictions

Jurisdiction	Relevant legislation	Description of change	Relevant legislation or wording
Tasmania, Australia	<i>Family Violence Act 2004</i>	Established the offences of economic abuse and emotional abuse or intimidation	The Act defines emotional abuse or intimidation as a course of conduct that a person knows, or ought to know, is “likely to have the effect of unreasonably controlling or intimidating or causing mental harm, apprehension or fear” in their spouse or partner (s 9(1)). This includes “limiting the freedom of movement of a person’s spouse or partner by means of threats or intimidation” (s 9(2)). Economic abuse is defined under s 8.
England and Wales	<i>Serious Crime Act 2015</i>	Established the offence of controlling or coercive behaviour in an intimate or family relationship (s 76).	An offence is committed if: <ul style="list-style-type: none"> • a person repeatedly or continuously engages in controlling or coercive behaviour towards another person; • the people are personally connected; • their behaviour has a serious effect on the other person; and • the person knows or ought to know that their behaviour will have a serious effect on the other person (s 76).
	<i>Domestic Abuse Act 2021</i>	Created a statutory definition of domestic abuse, emphasising that domestic abuse is not just physical violence, but can also be emotional, controlling or coercive, and economic abuse. Also extended the controlling or coercive behaviour offence to cover post-separation abuse.	The Act defines behaviour of a person (“A”) towards another person (“B”) is “domestic abuse” if: <ul style="list-style-type: none"> • A and B are each aged 16 or over and are personally connected to each other, and the behaviour is abusive (s 1(2)). Behaviour is “abusive” if it consists of any of the following— <ul style="list-style-type: none"> • physical or sexual abuse; • violent or threatening behaviour; • controlling or coercive behaviour; • economic abuse (see subsection (4)); • psychological, emotional or other abuse; and it does not matter whether the behaviour consists of a single incident or a course of conduct (s 1(3)).

6 One specific concern is an increase in the phenomenon of ‘dual arrest’ where the police officer arrests both individuals in a domestic violence event as both individuals meet the criteria for the offence (DeLeon et al., 2006). This may in turn decrease the ability of victims to involve the criminal justice system and/or increase the rate of women who are incarcerated for domestic violence, particularly for those not perceived as ‘good victims’ (Walklate & Fitz-Gibbon, 2021; Watego et al., 2021).

7 There have also been several closely related pieces of legislation, such as the reform to the French penal code recognising psychological violence in 2010 and a bill introduced by the French in 2020 aimed at better protecting victims of domestic violence which referred to psychological violence between partners. A bill has also been introduced in the Connecticut and New York legislatures in the U.S., and California has passed legislation to allow coercive control to be used as evidence of domestic violence in family court proceedings. Evidence of coercive control has also been used in expert testimony as part of a partial defence for murder (Sheehy, 2018).

Table 1. Summary of coercive control legislation across jurisdictions (continued)

Jurisdiction	Relevant legislation	Description of change	Relevant legislation or wording
Scotland	<i>Domestic Abuse (Scotland) Act 2018</i>	Established the offence of abusive behaviour towards a partner or ex-partner (s 1).	<p>An offence is committed if:</p> <ul style="list-style-type: none"> • a person engages in a course of abusive behaviour toward their partner or ex-partner (s 1(1)(a)); • a reasonable person would consider the course of behaviour to be likely to cause the partner to suffer physical or psychological harm (s 1(2)(a)); and • that the person either intends to cause physical or psychological harm; or is reckless as to whether their behaviour causes physical or psychological harm (s 1(2)(b)). <p>Section 2(2)(a) of the Act specifies that violent, threatening or intimidating behaviour directed toward a partner or ex-partner constitutes abusive behaviour.</p> <p>Sections 2(2)(b) and s 2(3) broaden the definition of abuse to include behaviour that produces ‘relevant effects’ such as:</p> <ul style="list-style-type: none"> • making their partner or ex-partner dependent on or subordinate to them; • isolating their partner or ex-partner from friends and relatives; • controlling, regulating or monitoring their partner or ex-partner’s day-to-day activities; • restricting their partner or ex-partner’s freedom of action; or • frightening, humiliating, degrading or punishing their partner or ex-partner.
Ireland	<i>Domestic Violence Act 2018</i>	Established the offence of coercive control (s 39).	<p>An offence is committed if a person knowingly and persistently engages in behaviour that:</p> <ul style="list-style-type: none"> • is controlling or coercive; • has a serious effect on a relevant person; and • a reasonable person would consider likely to have a serious effect on a relevant person (s 39(1)).
Hawaii, United States of America	<i>HB2425 HD1 SD1: Relating to Domestic Abuse</i>	Amends the definition of “domestic abuse” under Hawaii’s insurance laws and laws relating to domestic abuse protective orders to include coercive control between family or household members.	<p>Domestic abuse was amended to include coercive control, which was defined as (§ 5):</p> <p>“Coercive control” means a pattern of threatening, humiliating, or intimidating actions, which may include assaults, or other abuse that is used to harm, punish, or frighten an individual. “Coercive control” includes a pattern of behaviour that seeks to take away the individual’s liberty or freedom and strip away the individual’s sense of self, including bodily integrity and human rights, whereby the “coercive control” is designed to make an individual dependent by isolating them from support, exploiting them, depriving them of independence, and regulating their everyday behaviour including:</p> <ol style="list-style-type: none"> 1. Isolating the individual from friends and family; 2. Controlling how much money is accessible to the individual and how it is spent; 3. Monitoring the individual’s activities, communications, and movements; 4. Name-calling, degradation, and demeaning the individual frequently; 5. Threatening to harm or kill the individual or a child or relative of the individual; 6. Threatening to publish information or make reports to the police or the authorities; 7. Damaging property or household goods; and 8. Forcing the individual to take part in criminal activity or child abuse

Table 1. Summary of coercive control legislation across jurisdictions (continued)

Jurisdiction	Relevant legislation	Description of change	Relevant legislation or wording
Northern Ireland	<i>Domestic Abuse and Civil Proceedings Act (Northern Ireland) 2021</i>	Established the offence of domestic abuse (s 1).	<p>The Domestic Abuse and Civil Proceedings Act (Northern Ireland) 2021 established the offence of domestic abuse (s 1). An offence is committed if a person ('A'):</p> <ul style="list-style-type: none"> engages in a course of behaviour that is abusive of another person ('B'); A and B are personally connected to each other at the time; a reasonable person would consider the behaviour to be likely to cause B to suffer physical or psychological harm; and that person A either intends to cause physical or psychological harm, or is reckless as to whether their behaviour causes physical or psychological harm (s 1). <p>Section 2 of the Act specifies that violent or threatening behaviour towards person B constitutes abuse, as does behaviour directed at B, their child or another person, that produces 'relevant effects'. These include:</p> <ul style="list-style-type: none"> making person B dependent on or subordinate to them; isolating person B from friends, relatives or other sources of social interaction and support; controlling, regulating or monitoring person B's day-to-day activities; depriving of, or restricting person B's freedom of action; or making person B feel frightened, humiliated, degraded, punished or intimidated (s 2).
New South Wales, Australia ⁸	<i>Crimes Legislation Amendment (Coercive Control) Bill 2022.</i>	Established the offence of coercive control	<p>Sch 1 of the Crimes Legislation Amendment (Coercive Control) Bill 2022 introduces a stand-alone offence for coercive control (s 54D(1), <i>Crimes Act 1900</i> (NSW)). The offence can only be charged if five elements are met:</p> <ol style="list-style-type: none"> An adult (18 years or older) must engage in a 'course of conduct'. This means engaging in behaviour repeatedly or continuously. That course of conduct must be 'abusive behaviour'. This means behaviour that involves violence, threats or intimidation and/or coercion or control of the person against who the behaviour is directed. A reasonable person must consider that the 'abusive behaviour' would, in all circumstances, be likely to cause the other person to fear that violence will be used against them, or have a serious impact on their day-to-day activities. The 'abusive behaviour' must be directed against a current or former intimate partner. The offender must intend their course of conduct to cause physical or mental harm, or the offender must have been reckless about whether this could happen.
Queensland, Australia	<i>The Domestic and Family Violence Protection (Combating Coercive Control) and Other Legislation Amendment Bill 2022</i>	Amends to address coercive control: <ul style="list-style-type: none"> Coroners Act 2003 Domestic and Family Violence Protection Act 2012 Evidence Act 1977 Oaths Act 1867 Penalties and Sentences Act 1992 Public Guardian Act 2014 Telecommunications Interception Act 2009 Youth Justice Act 1992 Criminal Code Act 1899 	<p>Large range of changes. For example, s 31 amends the definition of domestic violence in the <i>Domestic and Family Violence Protection Act 2012</i> as:</p> <p>(1A) Behaviour, or a pattern of behaviour, mentioned in subsection (1)—</p> <ol style="list-style-type: none"> may occur over a period of time; and may be more than 1 act, or a series of acts, that when considered cumulatively is abusive, threatening, coercive or causes fear in a way mentioned in that subsection; and is to be considered in the context of the relationship between the first person and the second person as a whole.

Note: Information for this table was partially drawn from Otter and Bosanko (2022)

⁸ Hall (2021) provides an overview of the major changes to NSW criminal justice since 2010, noting that there has been a focus on domestic violence and sexual assault.

The current study

The aim of this study was to construct a measure of coercive control from the free-text domestic violence narratives recorded by NSW police. Our measure is constructed by identifying specific behaviours (coercive control behaviours) from the narratives using a set of rules (and related dictionaries) to define which *snippets* of text would constitute a mention of coercive control behaviour.⁹ As there were too many police narratives (526,787 narratives in total) to annotate manually, we employed a text mining method to identify these behaviours. We limited our study to detecting behaviours that are not well captured by the NSW police Force incident categories, rather than detecting behaviours such as physical and sexual abuse which are already well defined.

The free-text police narratives capture what happened at an event from the police's perspective and do not aim to ascribe motivation or identify patterns in behaviour. This may result in significantly biased estimates or underestimates of the prevalence of coercive control behaviours at these events, as they will depend on what the police have chosen to include in a narrative that was not intended to be used for this purpose. We have used this data source as it is one of the only sources of rich retrospective information on behaviours during and preceding domestic violence events for this period in NSW. Despite the significant limitations of this data, the resulting measure may still be useful as a supplementary source of information for both descriptive and predictive purposes.

We investigated two specific ways that this measure could be practically useful:

1. First, we examined how often our measures of coercive control behaviour agree with existing measures of the same behaviour, where these exist. The existing measures are drawn from the offence categories captured by NSW police as fixed fields in COPS. This provides some indication of the reliability of our measures for behaviours where no other information exists, and how a text mining measure could expand the set of events where these behaviours are detected.
2. Second, we compared the predictive power of a gradient-boosted regression model, with and without the inclusion of our coercive control measures, in predicting violent domestic violence offences within 12 months of the index event by the same person of interest. This allows us to assess the marginal benefit of constructing these variables for the purpose of prediction, identifying which offenders are at highest risk of reoffending.

METHOD

Data source

Two data sources were used in this study.

1. Information on domestic violence events captured in the NSW Police Force's Computerised Operational Policing System (COPS). This includes:
 - the structured data that is captured in the database (e.g., demographic information, the offence category for the incident, location and time of the offence) in the COPS fixed fields and;
 - the free-text case notes (known as police narratives).
2. Data on the offending history of each person of interest (POI) identified in these domestic violence events, extracted from the Reoffending Database held by the NSW Bureau of Crime Statistics and Research (BOCSAR). The Reoffending Database links all finalised NSW criminal court appearances since 1994 for an individual and relevant demographic information for those individuals.

⁹ As coercive control was not an offence when the narratives were recorded, we are unable to 'learn' a definition of coercive control from the data as there is no labelled training set.

We also constructed a third dataset by applying the text mining system to the police narratives in order to extract various coercive control behaviours. Each row in this dataset is an identified coercive control behaviour and includes the category and subcategory of high level abuse types that the behaviour belongs to (i.e., physical abuse, psychological abuse, stalking, harassment), as well as the snippet of text that indicated this behaviour. This categorisation is explained in detail in the method section below. Each narrative can have multiple instances of behaviours from the same category and/or subcategory of coercive control behaviour as well as multiple instances of behaviours falling within different categories.¹⁰

Sample

1. Domestic violence data from the NSW Police Force's Computerised Operational Policing System (COPS)

The sample for this study was selected by identifying all events that occurred between 1 January 2009 and 31 March 2020 that included domestic violence (DV) flagged incidents under the *Crimes (Domestic and Personal Violence) Act 2007* (NSW) and then requesting the narratives associated with these events from NSW police. Non-criminal incidents (e.g., those recorded as *domestic violence – no offence*) were excluded from the sample. This data does *not* have a variable that identifies whether coercive control was present, as coercive control was not an offence when this data was collected.

The final sample consisted of a total of 526,787 event narratives, with each narrative associated with one or more criminal incidents in the COPS system. The sample included information for 867,220 domestic violence incidents across 223,645 unique primary persons of interest.¹¹ For each of these events, we have demographic information for both the person of interest (POI) and the victim(s), along with data on the type of premises where the event occurred, whether the POI was proceeded against in relation to the event and the offence categorisation recorded by the NSW Police Force.

It is important to note that this sample contains all DV related criminal events recorded by police, not just those that relate to intimate partner relationships (where the concept of coercive control is arguably more meaningful). We included all events as information on the relationship between the victim and the POI was not available for all offence categories.¹² In our sample there were 468,408 (51%) incidents where we could not establish whether the incident occurred within the context of an intimate partner relationship.¹³ This limitation of the data means that restricting the sample to just those in intimate partner relationships would also limit the offences and behaviours that are captured in our analysis, and potentially impact our understanding of population prevalence and/or predictive power.

2. Offending data from NSW BOCSAR Reoffending Database (ROD)

We constructed our second dataset by selecting the relevant predictors and outcome variables from the NSW BOCSAR Reoffending Database for the primary POIs who were identified in relation to the domestic violence events extracted from COPS (described above). Of the 526,787 domestic violence events in our sample, we were able to link POIs from 497,739 (94.5%) events with our Reoffending Database.

3. Text mining data derived from police narratives

This dataset was constructed by applying the text mining system to each of the narratives extracted from COPS. We identified 852,162 occurrences of coercive control behaviours across the 526,787 events, with 299,689 (56.8%) narratives having at least one coercive control behaviour identified in the text.

¹⁰ As our text mining system captures 'snippets' that identify the behaviour, if a single behaviour is described multiple times in a narrative we identify it multiple times. For this reason, we consider the behaviour being identified at least once in the narrative as being more reliable than the count of behaviours in total.

¹¹ For each event, there can be multiple POIs attached to the incident. We have used the primary POI for each incident that we linked to a Master Personal Detail Identifier (MSPDI) in ROD.

¹² Relationship information was only available for 6 out of 49 offence categories in our sample - Homicide, Assault, Sexual offence – assault, Sexual offence – other, Offence against the person other, Robbery.

¹³ The most common incident categories that we do not have a relationship type for are Breach of AVO (16.9%), Malicious damage (12.7%), Miscellaneous (4.7%), Resist/hinder/assault officer (1.5%) and Firearms legislation (1.4%).

Identifying coercive control from police narratives

In this paper, we construct and then assess a measure of coercive control behaviours that we have constructed from the police narratives using a dictionary and rule-based text mining approach. This approach allows us to apply a sophisticated criterion to identify 48 different subcategories of coercive control behaviours consistently across our corpus of 526,787 police narratives – something that would have been impractical using manual annotation.

We use the text mining system to detect behaviours from the police narratives that could be present in a relationship typified by coercive control (i.e. coercive control behaviours), using rules that define and capture common syntactical patterns. Our measure is then defined as the number of times each type of behaviour is detected in the event narrative, keeping in mind that our observations are limited to behaviours observed and recorded by the police at each event.

Defining coercive control behaviours

Coercive control in domestic violence is difficult to define and codify for a text mining approach, as it describes *any* set of tactics aimed at dominating and controlling another individual within a relationship, rather than a defined set of behaviours. Furthermore, a behaviour that is used to coercively control an individual in one relationship may not necessarily constitute coercive control in another relationship. Many authors (Johnson, 1995; Stark, 2007) have argued that this means we cannot be confident that coercive control is taking place if a behaviour (or set of behaviours) is present (or conversely, that coercive control is not taking place if a specific set of behaviours is not present). While this implies that any source of data could be an incomplete observation of coercive control, this should not preclude attempts to improve the accuracy of the data which is available through collecting richer information on the behaviours taking place and refining how these behaviours are interpreted.

One potential source of information within existing administrative data collections are police narratives of domestic violence events. These are free text descriptions recorded by police, which capture a much richer description of what occurred at the event than is coded by police in incident data. Police narratives aim to capture what happened at an event from the police's perspective (and recording requirements) and generally do not aim to ascribe motivation or identify patterns in behaviour. As such, they are unsuited to detecting coercive control using a definition which is based on the intention of the abuser and/or the impact on the victim. As police narratives could produce an underestimate and/or biased estimate of prevalence due to the limitations of our method and the recording of coercive control in the police narratives, we have provided analyses where our measure could supplement and improve on our existing measures rather than being used uncritically as a measure of prevalence.

We aimed to capture a non-exhaustive set of coercive control behaviours as identified by Stark (2007). We exclude physical and sexual assault which, although potentially can be a part of coercive control behaviour, are likely to be captured accurately in the existing offence categorisations used by police. The text mining system does aim to detect intimidation, property damage and theft, despite existing police incident categories for these offences, as the narratives may identify additional occurrences of these behaviours within events.

Text mining approach

Our approach uses a combination of rules based on common syntactical patterns and dictionaries for the recognition of various abuse types. For example, the phrase “refused to let him in the house” would be detected as “premise lockout”, as would sentences that are similar such as “refused to let her in the house” or “refused to let him in the home”. Our measure is defined as the number of times each type of behaviour is detected in the event narrative, with observations limited to behaviours observed and recorded by the police at each event.

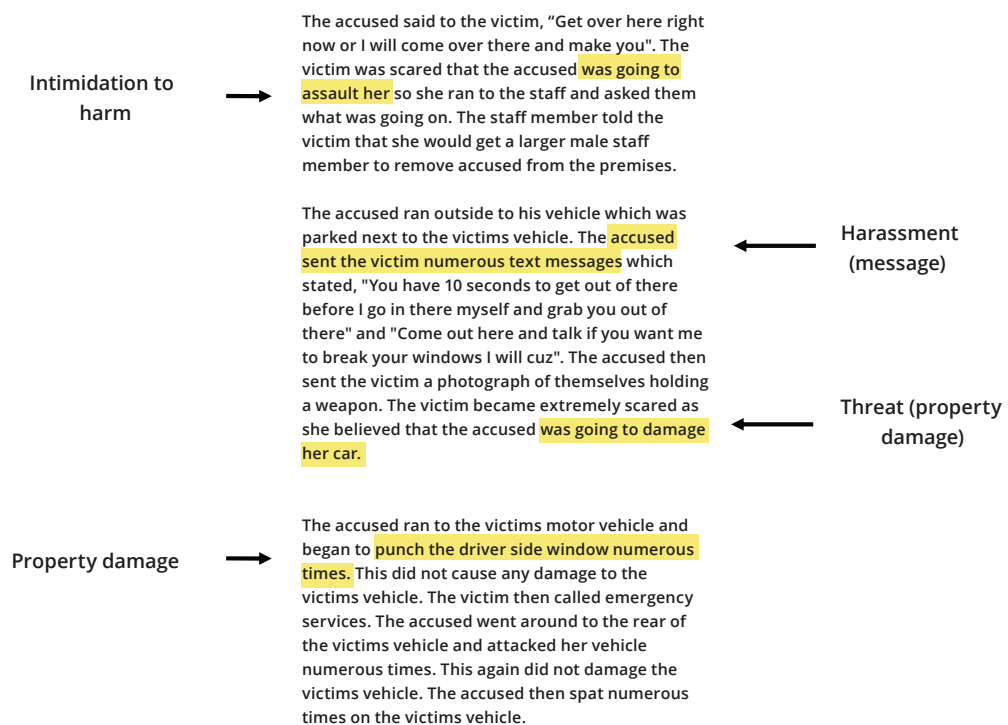
An existing text mining methodology was used to extract coercive control related behaviours in the domestic violence narratives. This method has already been evaluated and applied to a previous set of domestic violence event narratives for the automatic extraction of abuse types (Cunningham et al., 2013; Karystianis et al., 2019). It was implemented through the General Architecture for Text Engineering (GATE), which is a family of open-source text analysis tools and processes. Since the aim of this study is to capture behaviours related to coercive control, we restructured the previous method to focus solely on the identification of non-physical behaviours.

Our method included the following steps:

1. Designing rules that work in conjunction with the dictionaries to identify common sentence types and syntaxes for each coercive control behaviour;
2. Expanding dictionaries to capture a wider range of (valid) possibilities for each behaviour by allowing a range of pronouns, nouns and possible verbs for each syntax 'rule' that would still leave the *snippet* a valid description of a coercive control behaviour;
3. Mining the text to see which *snippets* match each of the decision rules, with each rule categorised under a specific coercive control behaviour;
4. Aggregating this information to the event level, to identify which behaviours happened at what times for which individuals.

Figure 1 below provides an example of a narrative with extracted coercive control related behaviours and their respective categories.

Figure 1. Example of marked up narrative



Dictionaries

We used fifteen existing dictionaries of non-physical abuse types which were manually crafted by Karystianis et al. (2019) for the previous project as a starting point for our method. We expanded these dictionaries by incorporating additional terms that describe non-physical abuse types after manually inspecting and annotating 200 randomly selected narratives from the current dataset. The “threat” dictionary from the previous study was subsequently broken down and expanded into twelve additional dictionaries that contained specific threat types towards the victim. These were then supplemented by an additional set of 11 dictionaries comprised from terms related to non-physical abuse types, such as various forms of messages that can be used to harass a victim and social media types. The dictionaries were manually assessed by two experts in forensic psychiatry and stalking¹⁴ to improve their accuracy.

Rules

We developed rules to identify coercive control behaviour using recurring syntactical patterns (e.g., “X sent over Y text messages to Z’s phone”). This allows us to define many variations of a phrase that identifies a specific coercive control behaviour. The syntactical patterns utilise:

- frozen syntactical expressions as anchors for certain elements built through specific verbs;
- noun phrases, and prepositions (e.g., “proceed to stalk the victim”) and;
- semantic placeholders identifiable through the application of the manually crafted dictionaries (e.g., all possible synonyms describing a POI, such as “POI,” “perpetrator,” and “offender”).

Once we had developed a viable set of rules and dictionaries, we looked at the prevalence of each behaviour, as well as the conceptual coherence of each set of rules and dictionaries to redefine our categories and subcategories. We went through three major iterations of this text mining system, each time making substantive changes to the coercive control behaviours that were being captured and how they were defined at each step.

For this study, we eventually settled on 618 rules and 36 dictionaries to identify 48 coercive control behaviours in six categories. These are presented below, with examples of snippets that were captured in each category.

Table 2. Categories and subcategories of coercive control behaviours

Category	Subcategory	Snippet
Controlling behaviour	Blackmail	threatened the victim with exposing the relationship to her family
	Controlling behaviour (other)	been a history of controlling behaviour
	Control over employment	not allow her to work
	Control over finances	taking the money
	Home isolation	restraining the victim in the home
	Internet search	Police also located numerous google searches with the victim’s name
	Isolation from children	not let her see her son
	Isolation from family	not letting her see her family
	Isolation from friends	not allowing her to socialise
	Non-consensual sharing of intimate material	POI is using this video as leverage
	Premise lock out	refused to let him in the house
	Privacy violation	pn2 went through her mobile phone
	Stated surveillance	“I have put the house under 24 hour video and audio surveillance”
	Use of tracking device	placing the tracking device on the victim’s vehicle
	Unauthorized information dissemination	recorded her without her permission
Deprivation of basic necessities	NA – no narratives matched this subcategory ¹⁵	

¹⁴ Dr Sunny Wade and Dr Stephen Allnutt.

¹⁵ We did not detect any behaviours in this category over our corpus.

Table 2. Categories and subcategories of coercive control behaviours (continued)

Category	Subcategory	Snippet
Harassment	Call	Threatening phone calls
	Email	received several emails from the accused
	Letter	and leaving love letters
	Location	asked the accused to leave the house several times
	Message	received no less than 25 text messages
	Obsession	is obsessed with her
	Social media	received 404 Facebook messages
	Stalking via car	defendant drove past the PINOP's address several times
	Stalking at location	accused has parked his vehicle across the road from the victim's residence
	Stalking (unspecified)	charged with stalk
	Trespassing	forced his way into
	Unwanted gift	had given her as gifts
Unspecified	harassed by the accused	
Intimidation and threats	Intimidation to harm	was going to hit him
	Threat to kill (stated)	"I'll kill the c****"
	Threat to sexual assault (stated)	"I will f*** you"
	Intimidation (unspecified)	and making threats
	Intimidation with an object	held it to his chest
	Threat to property damage (stated)	"I'm going to smash the door"
	Threat to harm/kill animals (stated)	"I'll kill your dog"
	Threat to harm third person (stated)	"I will cut her"
	Veiled threats for harm (stated)	"I'm going to make your life hell"
	Threat to harm animals	previously threatened to kill the victim dogs
	Threat to harm others	INFT threatening to kill his partner
Threats taking children away	threatened to take the baby	
Self-harm and suicide (acts and threats)	Acts of self-harm	Young Person has began to cut her hand
	Self-harm threat	Accused has made threats of self-harm
	Suicide threat (stated)	"I am going to kill myself"
Property damage and theft	Arson	POI set fire
	Property damage	and malicious damage
	Withholding personal effects	accused has snatched her mobile phone
Verbal abuse	Verbal abuse	and yell to the victim

Evaluation

The methodology outlined above has previously been evaluated against a set of 100 randomly selected domestic violence narratives with promising performance at the event narrative level in identifying various abuse types (Karystianis, 2019). We calculated precision (90.2%), recall (89.6%) and F1-Score (89.8%) respectively from this evaluation, which suggest reliable performance. As we did not include additional abuse types, a new evaluation at this stage was deemed unnecessary. However, to ensure consistency in performance, 100 additional narratives classified with the offence "stalking/harassment" were randomly selected from the current dataset and were manually inspected to ensure that all types of non-physical abuse were captured by the method and that there were no false positives.

Comparing the measure of coercive control behaviour to existing categories

Once we applied the text mining system on the corpus of police narratives, we compared our measure of coercive control to existing measures of the behaviour that were already being recorded through fixed fields in COPS. This analysis helped us understand how our measure of coercive control differed from what is already being captured and utilised by the NSW police.

We conducted this exercise for the following 5 existing police incident categories:

- **Stalking:** We compared this incident type with the stalking categories derived from our measure (stalking via car, stalking at location, stalking unidentified).
- **Intimidation:** We compared the incidents categorised as “bullying/harassment or intimidation” or “intimidation” with the intimidation categories derived from our measure (intimidation to harm, intimidation to kill stated threats, intimidation to sexual assault stated threat, intimidation unspecified, intimidation with object).
- **Use carriage service to menace/harass/offend:** We compared the incidents categorised as “indecent communication” or “telecommunications offence” with the harassment categories (i.e., call, message, email, letter, social media).
- **Property damage:** We compared the incidents categorised as “malicious damage to property¹⁶” with the derived subcategory “property damage”.
- **Trespass:** We compared the incidents categorised as “trespass with firearm” and “trespass” with the subcategory derived “trespassing” subcategory.

For each of these comparisons, we showed a simple confusion matrix that split the sample into the events where the incident category and the text mining measure both categorised the event as TRUE or FALSE and where the text mining measure detected the behaviour but the police did not use the corresponding incident category or vice versa. We would not expect the incident category and the text mining measure to align exactly, as police may not always record the necessary detail or sequence of events in the narrative that is required for the text mining process to identify the behaviour, nor does the presence of the behaviour in the narrative imply that it met the criteria for the police incident category.

Using the measure of coercive control behaviour to predict future violent offending

Finally, we assess how useful our measure of coercive control is as a predictor of future violent offending in domestic relationships.¹⁷ To do this, we build a predictive model using a gradient boosted decision tree approach (implemented using the xgboost package in R; Chen & Guestrin, 2016) and report standard metrics of predictive accuracy (Accuracy Area Under the Curve (AUC), Precision, Sensitivity¹⁸) with and without our measures included as predictors. We estimate this model after transforming the data to the POI / Event level, which means we have 479,984 observations of DV related POI / event pairs.¹⁹

We have chosen to use a gradient boosted model for two key reasons. First, this approach allows us to flexibly utilise a large set of variables for prediction while imposing minimal restrictions on model structure.²⁰ This means that the predictive accuracy is less dependent on how we have specified the

¹⁶ There are two additional subcategories that could be relevant, “public place - damage fountain/wall etc” and “public place - damage shrine/monument”, which have not been included. There are only 11 events in these subcategories combined, making them too small to materially impact the analysis.

¹⁷ For our definition of violent offending, we have used all offences that are in ANZSOC category 1, 2, 3, 5 or 6.

¹⁸ Briefly, the AUC can be interpreted as the probability that the model can discriminate between a randomly chosen positive and a randomly chosen negative observation, the precision is the proportion of positives detected by the model that are actually positives, and the sensitivity is the proportion of positive values detected by the model.

¹⁹ This differs from the total number of events (526,787) in our data for two reasons – first, an event can have multiple POIs. Secondly, we have filtered out all events that do not have a POI identified as their outcome is missing.

²⁰ For this analysis, we have left all tuning parameters at their default values, as we are primarily interested in the marginal value of our measure of coercive control rather than optimising the effectiveness of the predictive model.

model (e.g. which variables are used and how) when compared with a more standard linear regression approach. Second, gradient boosted models implemented with xgboost have demonstrated excellent performance compared to other comparable machine learning algorithms across a range of prediction problems and settings (Chen & Guestrin, 2016).

We first split the data by date into a training set (80% of the data) and the test set (20% of the data)²¹, and then further split the training data to tune the parameters of the model using 5-fold Monte Carlo cross-validation.²² We conduct this process five times for each model and conduct an efficient grid search across 7 parameters using an ANOVA race approach. (Kuhn 2014, implemented via the finetune package in R). We use the F-score as our tuning metric as our data is significantly imbalanced²³, with very few 'positive' cases being detected if accuracy is used as the tuning metric.

Once we have tuned our parameters, we trained each model on the full training set and evaluated the predictions (e.g. calculated the metrics of predictive accuracy presented) based on the remaining 20%. We tuned and estimated three models: (1) a model with just the variables that are routinely used to predict reoffending; (2) a model with just our measures of coercive control derived from the text mining process; and (3) a model with both sets of variables. The full set of variables that we used for prediction are provided in the Variables section of this report. The analysis was conducted at the incident level, which allows us to differentiate between different POIs who commit different offences within the same event.

We present the metrics of predictive accuracy for all three models to provide a "hard test" of the predictive value of our measures over and above the readily available variables typically used to predict reoffending (e.g., age, gender, prior offending, prior imprisonment). Relying on simple "importance" measures or coefficient characteristics may significantly overstate how useful a variable is for the prediction task, as it is common for machine learning approaches to have similar predictive power on a prediction task while using completely different sets of predictors (Mullainathan & Spiess, 2017). Comparing metrics across separate models allows us to better assess the marginal improvement in prediction caused by including our measures of coercive control.

Variables

Text mining data derived from event narratives

For each event narrative in our dataset, we constructed a measure of coercive control behaviours using our text mining method. This allows us to identify how many of the 48 subcategories of coercive control behaviour are detected for each event, as described in the previous section. In our sample, approximately one in five events had multiple occurrences of the same coercive control behaviour.²⁴

Variables used for predicting future violent reoffending

The outcome variable for our predictive model was whether the Person Of Interest (POI) identified in the event committed a new domestic violence related *violent* offence within 12 months of the index event. A violent offence is defined as any offence falling within the following Australian and New Zealand Standard Offence Classification (ANZSOC) categories:

- 01 Homicide and related offences
- 02 Acts intended to cause injury
- 03 Sexual assault and related offences
- 05 Abduction, Harassment, and other offences against the person
- 06 Robbery, extortion, and related offences

21 E.g. our data was trained on all events 1 January 2009 to 25th March 2018, and then evaluated for accuracy on data from 26th March 2018 to 31st of March 2020.

22 This means that we randomly selected 75% of the training data and use the other 25% of the training data to evaluate the tuning parameters.

23 6.7% cases are followed within 12 months with DV related violence by the same POI.

24 This is dominated by multiple occurrences of property damage, verbal abuse, and intimidation within the event, with these making up 90% of these events.

We then use two separate sets of predictors, a “standard set” that is already available and typically used for predicting reoffending, and our text mining variables. The standard set of demographic and prior offending variables that were used are described below.

- Age: An integer that indicates the age of the POI at the index event
- Gender: A categorical variable that indicates the gender of the POI, recorded as “Male,” “Female” and “Unknown”
- Aboriginal status: A logical variable which is TRUE if the POI was ever recorded as being Aboriginal, FALSE otherwise
- Socio-economic Index for Areas (SEIFA): A categorical variable with five values, one for each quartile of the SEIFA distribution and one for missing values. The higher the score, the greater the advantage of the area where the individual resides
- Alcohol related: A logical variable which is TRUE if the incident was flagged as alcohol related, FALSE otherwise
- DV related: A logical variable which is TRUE if the incident was flagged as domestic violence related, FALSE otherwise
- Incident category Level 1: Category of incident recorded by police, factor variable with 49 categories
- Incident category Level 2: Sub-category of incident recorded by police, factor variable with 216 categories
- Victim’s relationship to POI: Description of the victim’s relationship to the POI, factor variable with 13 categories (and missing)
- Court appearances: An integer variable that indicates the number of court appearances by the POI in the five years before the index event
- Charges: An integer variable that indicates the number of charges against the POI in the five years before the index event
- Court appearances where POI found guilty: An integer variable that indicates the number of court appearances in the five years before the index event where the POI was found guilty
- Charges where POI was found guilty: An integer variable that indicates the number of charges in the five years before the event where the POI was found guilty
- DV related court appearances where POI found guilty: An integer variable that indicates the number of DV related court appearances in the five years before the event where the POI was found guilty
- DV related charges where POI was found guilty: An integer variable that indicates the number of DV related charges in the five years before the event where the POI was found guilty
- Court appearances where POI found guilty of a violent offence: An integer variable that indicates the number of court appearances in the five years before the event where the POI was found guilty of a violent offence
- Charges where POI was found guilty: An integer variable that indicates the number of charges in the five years before the event where the POI was found guilty

The text mining variables that we have used are presented in Table 2, as well as in the results section. We have included a variable for each of the six categories and each of the 48 subcategories, capturing the number of snippets which match each of the categories or subcategories for each narrative. We have also aggregated these variables within event to create two additional variables that give the sum of the number of distinct categories/subcategories within each event. We then created an additional 56 variables (6 categories, 48 subcategories, 2 totals) that capture the number of matches for each of these categories/subcategories in the year prior to the event for the POI. These variables have been constructed to better capture coercive control as a ‘course of conduct’ rather than a behaviour observed at a single event.

Descriptive statistics

Below we present summary statistics for the variables used in our predictive modelling at the incident and POI levels, respectively. Summary statistics for the text mining variables are presented in detail in our results section.

Incident characteristics

Table 3 shows summary statistics for each incident in our extract of the COPS data, which includes all domestic violence flagged events from 1 January 2009 to 31 March 2020.

Table 3. Characteristics of domestic violence incidents from COPS dataset

Characteristic	N = 867,220 ¹
Police incident category	
Assault	330,673 (38%)
Breach Apprehended Violence Order	146,881 (17%)
Drug detection	8,091 (1%)
Firearms legislation	11,763 (1%)
Intention offence	5,930 (1%)
Judicial offences	30,934 (4%)
Malicious damage	109,730 (13%)
Miscellaneous	40,741 (5%)
Offence against the person, other	111,821 (13%)
Powers	6,445 (1%)
Resist/hinder/assault officer	13,058 (2%)
Sexual offence	20,801 (2%)
Traffic	6,731 (1%)
Other	23,621 (3%)
Gender of victim	
Female	557,785 (64%)
Male	204,203 (24%)
Unknown	105,232 (12%)
Age of victim	
Median (IQR)	34 (24, 44)
N missing (% missing)	105,277 (12%)
Gender of POI	
Female	157,452 (18%)
Male	688,896 (79%)
Unknown	20,872 (2%)
Age of POI	
Median (IQR)	32 (24, 41)
N missing (% missing)	22,958 (2.6%)
Relationship of victim to the POI	
Intimate partner	
Boy/Girlfriend (Incl Ex-Boy/Girlfriend)	95,838 (11%)
Spouse/Partner	90,362 (10%)
Ex-Spouse/Ex-Partner	68,713 (8%)
Parent, Guardian, or other Authority figure	
Parent/Guardian of Victim	37,108 (4%)
Carer	4,479 (1%)
Person In Authority	643 (0%)
Family	
Member Of Family - Other	39,893 (5%)
Child (Incl Step/Foster Child) Of Victim	34,361 (4%)
Sibling	24,467 (3%)
Other	
Other Known Person- No Relationship	23,380 (3%)
Household Member (Incl Former Household)	17,460 (2%)
Not Known to Victim	5,203 (1%)
Unknown	425,313 (49%)

¹ n (%)

The most common type of police incident category for events included in the Sample was assault (38%), but apprehended violence order breaches (17%), malicious damage (13%) and other offences against the person (e.g., bullying, harassment, intimidation, abduction, stalking) (13%) were also common. Women made up at least 64% of the victims involved in these incidents, while 79% of POIs were men. The median age of victims was 34, which is slightly older than the median age of POIs (32). As noted in the sample section, information on the relationship between the victim and the POI was missing for 49% of incidents. The majority of events where this information was available involved violence between intimate partners (including incidents involving Boy/Girlfriends (11%), Spouse/Partners (10%), and Ex-Spouse/Ex-partners (8%)).


RESULTS

Estimated prevalence of coercive control behaviour

In this section, we present the prevalence of coercive control behaviours in domestic violence related events using the measure constructed by our text mining system. As our text mining measure aimed to capture coercive control behaviours that were not already well captured by police, our estimates omit coercive control behaviours such as physical and sexual assault from our estimates. In the following section, we assess how well our measure performs when compared with the offence category captured by police.

Overall, there were 852,162 behaviours captured across 526,787 events, with 57% of events having at least one coercive control behaviour detected. These events were associated with 223,645 unique persons of interest. We detect three or more unique subcategories of coercive control behaviour²⁵ in 8% of events.²⁶ Figure 2 shows the number of unique subcategories detected in each event.

Figure 2. Number of coercive control subcategories detected in each event




²⁵ We present unique subcategories rather than coercive control behaviours detected here as we are unable to distinguish between incidence of behaviours and the number of times the same behaviour is mentioned.

²⁶ These estimates are difficult to compare with those presented in Boxall (2021) as they are constructed from very different samples, look at different behaviours, use different methodologies to capture them and use different time periods for the behaviour to occur.

Figure 3 presents the number and relative proportion of narratives where at least one instance of the six main categories of coercive control behaviours were detected. As seen here, property damage and theft (26%), intimidation and threats (24%) and verbal abuse (23%) were the three most common categories of coercive control identified by our text mining system, with 53% of all events having a coercive control behaviour from at least one of these categories. In contrast, harassment was identified in only 6% of narratives, and behaviours falling in the “Self harm and suicide” or “Controlling behaviour” categories were detected in less than 3% and 1% of narratives respectively. In the next sections we break down three of these six categories (Intimidation and threats, harassment, self-harm and suicide and controlling behaviour) into their underlying subcategories. The other three categories (property damage and theft, verbal abuse, self-harm and suicide) are not broken down since they are essentially made up of one or two large subcategories.²⁷

Figure 3. Number of coercive control behaviours, by category




Intimidation and threats by subcategory

Figure 4 below presents all subcategories that make up the intimidation and threats category, along with the number of narratives where these behaviours were detected. The two most common subcategories of intimidation and threats identified through the text mining process were the subcategories of “intimidation unspecified” and “intimidation to harm.”²⁸ These two intimidation behaviours were present in over 105,000 events, or around 20% of all events in our sample. “Threat to kill” was the next highest category (6%), followed by “intimidation with object”, “threat to property damage”, and “veiled threats for harm”.

²⁷ For the property damage category, the subcategory “property damage” makes up 89% of observed events, “withholding of personal effects” making up 11% and the rest being made up by “arson”. The subcategory “verbal abuse” makes up 100% of the Verbal abuse category. Finally for self-harm and suicide, threats to commit suicide and threats to self-harm make up 51% and 48% of the observations in this category respectively, with the remainder being acts of self-harm.

²⁸ ‘Intimidation unspecified’ is a subcategory that captures mentions of intimidation or threats that have no other information provided (e.g. ‘had made threats’). ‘Intimidation to harm’ captures narratives that have mentions of committing physical violence (e.g. ‘was going to hit her’)


Figure 4. Number of intimidation and threats detected, by subcategory



Harassment by subcategory

Figure 5 presents all subcategories that make up the harassment category, along with the number of narratives where these behaviours were detected. The most common types of harassment detected by the text mining system were harassment via call (10,199 events), harassment via messaging (8,224 events) and stalking unspecified (7,151 events). These harassment types were detected in 1-2% of all DV related events. There were a further 1,663 events where trespassing was identified.


Figure 5. Number of harassment behaviours detected, by subcategory



Controlling behaviour by subcategory

Figure 6 presents all subcategories that make up the controlling behaviour category, along with the number of narratives where these behaviours were detected. Two things stand out from this figure. First, the vast majority of the detected behaviours fall within either the subcategory “controlling behaviour (other)” or “control over finances”. These behaviours were identified as being present in 1,312 and 1,239 DV events, respectively.

Figure 6. Number of controlling behaviours detected, by subcategory



Second, very few events contained behaviours from 10 of the 15 subcategories that make up the ‘controlling behaviour’ category. For example, “stated surveillance”, “non-consensual sharing of intimate material”, “blackmail” and “internet search” were present in just one DV event each. This could indicate a low underlying prevalence of these behaviours in police recorded incidents of domestic violence. Alternatively, it could reflect the inability of our text mining method to detect instances of these behaviours in the police narratives and/or police not observing and/or recording when these behaviours take place.

Comparing text mining measures of coercive control with existing police incident categories

Four of the behaviours detected by our text mining system (stalking and intimidation, property damage, use of communication services to menace/harass/offend and trespass) also constitute offences in and of themselves, and thus were recorded by NSW police in our data. In this section, we have shown where our measure of coercive control derived from the text mining of narratives and the fixed field incident categories within COPS agree and disagree on whether a behaviour is present.

This comparison provides some indication of the reliability of our coercive control measures for behaviours where no other information is available. As our methodology aimed to minimise the number of false positives by manually crafting each of the rules used to detect the behaviour, we interpret the additional events detected by the text mining system as likely being a “true” match of the behaviour but the behaviour had not been categorised by the police into the corresponding incident category. Conversely, the cases where the incident category is TRUE and the text mining measure is FALSE are

likely to be cases where the text mining measure was not able to detect the behaviour despite it having occurred. As our measure is broader than the criteria required for an incident category to be recorded, our measure can be seen as a supplementary source of information on whether the behaviour occurred at all.

For each comparison, we present a confusion matrix, which shows where the measures agree and disagree on the presence of each behaviour across the 526,787 event narratives that make up our dataset.

Stalking and intimidation

First, we looked at how well our measures of stalking and harassment align with the same incident category in COPS.²⁹ Table 4 shows that of the 526,787 events in our data, our two measures agree for 418,063 (79%) of the events, with 56,605 (11%) being categorised as stalking and/or intimidation by both the incident category and the text mining system, and 361,458 (69%) being categorised as having neither of those behaviours.

Our text mining measure detects stalking and/or intimidation in an additional 63,171 (12%) events. This indicates that nearly half of the incidents picked up by the text mining measure are not classified by police as such using the existing incident categories. There were a further 45,553 (9%) incidents that were not picked up using the text mining methodology but were classified as stalking and intimidation using the police incident categories. Given that the design of the text mining system prioritised minimising false positives, this suggests that our measure of coercive control behaviour can improve estimates of the prevalence of stalking and intimidation behaviours in DV incidents reported to police.

Table 4. Stalking and intimidation - confusion matrix using police incident categories

		Text mining	
		TRUE	FALSE
Incident category	TRUE	56,605 (11%)	45,553 (9%)
	FALSE	63,171 (12%)	361,458 (69%)

Property damage

For property damage, our text mining measure and the existing police incident category produced very similar results (Table 5), with the two measures agreeing on 467,758 events or approximately 89% of the time. However, the text mining measure picks up an additional 42,962 (8%) events containing property damage compared with the existing incident category, and only fails to detect 16,067 (3%) events that have been classified as property damage by the police.

Table 5. Property damage - confusion matrix using police incident categories

		Text mining	
		TRUE	FALSE
Incident category	TRUE	86,204 (16%)	16,067 (3%)
	FALSE	42,962 (8%)	381,554 (72%)

²⁹ Worth noting here that 74.2% of these are categorised under “intimidation”, 24.8% under “bullying/harassment or intimidation” and 0.00006% (4 incidents) under stalking. In comparison, our text mining method detects 5,217 incidents with stalking behaviour.

Use of communication services to menace/harass/offend

Next, we turn to the use of communication services to menace/harass/offend.³⁰ Table 6 below shows that the text mining and incident categories agree most of the time, with 488,961 (91%) events categorised identically using the two measures. This is likely driven by the low incidence of these offences, with both measures detecting this offence in less than 7% of the DV events. Again, similar to stalking and intimidation offences, more than half the events picked up by the text mining measure (10,908 events) were not picked up by the existing police incident categories.

Table 6. Use carriage service to menace/harass/offend - confusion matrix using police incident categories

		Text mining	
		TRUE	FALSE
Incident category	TRUE	7,634 (1%)	26,918 (5%)
	FALSE	10,908 (2%)	481,327 (91%)

Trespass

Table 7 shows that the existing incident category and the text mining measure agree for almost all events (522,369 or 99% of events), but this is again almost entirely due to the very high proportion of incidents where trespassing is not present. Like with the use of communication service to menace/harass/offend, the majority of trespassing that is categorised by the existing incident category is not captured by the text mining measure (and vice versa). This suggests that our text mining measure is not able to detect trespass cases effectively as currently constructed.

Table 7. Trespass - confusion matrix using incident categories

		Text mining	
		TRUE	FALSE
Incident category	TRUE	187 (0.0%)	2,942 (0.6%)
	FALSE	1,476 (0.3%)	522,182 (99.1%)

Using text mining measures of coercive control behaviour to predict reoffending

In this section, we evaluated whether our measures of coercive control behaviours were useful for predicting whether the POI will commit a further *violent* domestic violence related offence within 12 months of the event. To do this, we constructed models with and without the text mining variables included and compared the predictive power of each model to estimate the marginal impact of including the text mining variables in the prediction.³¹ In our data, 6.7% of POIs go on to commit another violent domestic violence incident within 12 months of the index event, making the imbalance of the classes a significant issue for prediction.

30 Formally this is given as ‘use of carriage service’, where a carriage service is defined in the *Telecommunications Act 1997* (Cth) as ‘a service for carrying communications by means for guided and/or unguided electromagnetic energy (s 7)’ e.g., phone calls, text messages and social media communications.
 31 Our measures attempt to capture a wide range of coercive control behaviours, but do not aim to detect physical and sexual assault as they are already captured in the police incident categories. Thus, our results do not reflect the additive value of identifying whether physical or sexual assault was present in the event (despite being coercive control behaviours) as these are already captured in the ‘standard predictors’.

Table 8 presents the summary measures of predictive performance for each of our models. Overall, the inclusion of our text mining variables adds essentially no value to our predictive model relative to the model with standard predictors (e.g. between row 2 and 3 in Table 8 below). We see an improvement in the accuracy of the model by 0.02 percentage points and the models differ in the Area Under the Curve (AUC)³² measure by less than 0.001. Our coercive control measures also have essentially no impact on the precision and recall of the model (0.001 increase for precision and 0.008 decrease for recall).

Note that the model using all the text mining variables is still able to achieve an accuracy of 71% due to the relative inability to correctly predict when a further violent event will occur³³, detecting 73% of future violent events when compared with the other two models (i.e. the recall is 0.38, compared with 0.52 and 0.51). This is also reflected in the much lower AUC value, 0.57 compared with 0.83.

Table 8. Performance metrics of predictive models - incident level, all domestic violence events

	Accuracy	AUC	Precision	Recall
1 Text mining variables	71.0%	0.57	0.07	0.38
2 Standard predictors (e.g., POI demographics and priors)	84.6%	0.83	0.17	0.52
3 Both standard predictors and text mining predictors	84.8%	0.83	0.17	0.51

DISCUSSION

This study used a text mining method to identify coercive control related behaviours from the free-text descriptions in police narratives of 526,787 domestic violence related events occurring between 1 January 2009 and 21 March 2020. This allowed us to measure the prevalence of 48 separate behaviours which we grouped into six categories. The value of our work is two-fold. First, our study explores whether previously unutilised information within the police narratives can be extracted effectively through text mining to construct a measure of coercive control. For behaviours that have been defined with a high level of granularity (e.g. isolation from child, tracking device, threats to animal), these are (to our knowledge) the only measures that exist of the prevalence of these behaviours in domestic violence events in NSW across this period. Second, this study provides an indication of how useful this measure of coercive control is in conjunction with existing data, either by supplementing our understanding of what behaviours have taken place in addition to the existing incident categories, or as additional variables in a predictive model.

We found that in 57% of domestic violence events at least one coercive control behaviour was reported, with property damage and theft (26%), intimidation and threats (24%) and verbal abuse (23%) being the three most common behaviours. In 8% of events, three or more distinct subcategories of coercive control behaviours were detected. Other forms of coercive control, such as technological abuse, were detected but much less frequently.

We then assessed the extent to which our measures overlapped with existing measures in the police data system. For the four behaviours that were captured by both the police incident categories and the text mining system (stalking and intimidation, use carriage service to menace/harass/offend, property damage and trespass), we found substantial overlap with police offence classifications from fixed fields. However, we also found that for all four of the categories examined, we would detect many more additional behaviours if the text mining measure was used in conjunction with the incident category (between 30% and 60% more incidents of coercive control). This is consistent with the behaviour being present or reported to police but not necessarily resulting in legal action being taken.

³² The AUC measures the probability of a model distinguishing between a random positive and negative value selected from the sample.

³³ Any model can achieve an accuracy of 93.5% by categorising every incident as FALSE, due to the imbalanced nature of the dataset.

We also assessed how useful our text mining measures were for predicting violent reoffending by comparing predictive models with and without our measure of coercive control and assessing differences in predictive performance. We found that our coercive control measures provided no improvement in prediction over and above the demographic, offence characteristic and prior offending variables that are typically used to predict domestic violence reoffending. This is consistent with Karystianis et al. (2021) who also find a small improvement in prediction if text-mining variables are included (although these authors used different metrics of predictive power). However, our findings are also consistent with previous research showing the limited effectiveness of risk assessment tools in predicting more severe domestic violence offending (Dowling & Morgan, 2019; Leung & Trimboli, 2022; Svalin & Levander, 2020).

The key limitation of our study is that our measure of coercive control is derived from police narratives and is thus critically dependent on how police have recorded these behaviours. Our measure is better thought of as a measure of how NSW police have observed and recorded coercive control in the domestic violence events they attend, rather than as an accurate measure of the prevalence of coercive control across the population. This limitation, when combined with the limitations of our text mining method, may lead to an underestimate and/or a biased estimate of the prevalence of coercive behaviours. Our method requires a) the behaviour to take place when the police are present and/or be reported to police, b) the coercive control behaviours to be recorded in the police narrative *and* c) for the behaviour(s) to be recorded in such a way that it matches one of the rules in our text mining system. Although tools such as text mining and predictive modelling can help us to process and then operationalise large volumes of information that may otherwise be intractable, they cannot overcome issues inherent in the data.

The inherent difficulty in defining, operationalising, and capturing coercive control presents challenges for evaluating the impact of the new coercive control legislation in NSW. Our estimates suggest that more than half of DV events that come to the attention of the police have at least one coercive control related behaviour reported (although a much smaller proportion, only 8%, had multiple reports within the same event). It is therefore possible that a large majority of persons of interest identified in relation to DV events, could potentially be charged with this new offence. Given the difficulties observing this offence, the way in which police record and enforce the offence (and the way in which victims report the behaviour) will largely determine how often the offence is used (and when) and the subsequent impact.

Overreach of the reforms is a particular concern for more vulnerable groups. The increased enforcement of domestic violence stalking/intimidation offences, for example, was a major contributor to the growth in the Aboriginal prison population between 2001 and 2015 (Weatherburn & Ramsey, 2016), and there are concerns that this new offence could have a similar impact on Aboriginal people. Higher rates of coercive control offences for a group could reflect a higher incidence of coercive control behaviours, but could also simply be due to greater levels of surveillance and enforcement directed toward those individuals.

REFERENCES

- Aguilar Ruiz, R., & González-Calderón, M. J. (2020). Predictors of severe intimate partner violence among antisocial and family-only perpetrators: Victims' and offenders' characteristics. *Journal of Interpersonal Violence, 37*(5-6), NP2791-NP2822.
- Barlow, C. (2022). 'How can you capture what is hidden?' Police body-worn cameras and coercive control. *Journal of Gender-Based Violence, 7*(1), 1-15.
- Barlow, C., Johnson, K., Walklate, S., & Humphreys, L. (2020). Putting coercive control into practice: Problems and possibilities. *The British Journal of Criminology, 60*(1), 160-179.
- Brennan, I. R., Burton, V., Gormally, S., & O'Leary, N. (2019). Service provider difficulties in operationalizing coercive control. *Violence Against Women, 25*(6), 635-653.
- Boxall, H., & Morgan, A. (2021). *Experiences of coercive control among Australian women* (Statistical Bulletin no. 30). Canberra: Australian Institute of Criminology.
- Chen, T., & Guestrin, C. (2016). XGBoost: A Scalable Tree Boosting System. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 785-784*.
- Cunningham, H., Tablan, V., Roberts, A., & Bontcheva, K. (2013). Getting more out of biomedical documents with GATE's full lifecycle open source text analytics *PLoS computational biology, 9*(2), e1002854.
- DeLeon-Granados, W., Wells, W., & Binsbacher, R. (2006). Arresting developments: Trends in female arrests for domestic violence and proposed explanations. *Violence Against Women, 12*(4), 355-371.
- Dobash, R. E., & Dobash, R. (1979). *Violence against wives: A case against the patriarchy* (pp. 179-206). New York: Free Press.
- Domestic Violence Death Review Team (2020), Report 2017-2019, Retrieved on 18th January 2023 https://www.coroners.nsw.gov.au/content/dam/dcj/ctsd/coronerscourt/documents/reports/2017-2019_DVDRT_Report.pdf
- Donovan, C., & Hester, M. (2014). *Domestic violence and sexuality: What's love got to do with it?* Bristol: Policy Press. Retrieved from <https://sure.sunderland.ac.uk/id/eprint/8373/>
- Dowling, C. & Morgan, A. (2019). *Predicting repeat domestic violence: Improving police risk assessment*. (Trends & issues in crime and criminal justice no. 581). Canberra: Australian Institute of Criminology. Retrieved from <https://www.aic.gov.au/publications/tandi/tandi581>
- Frankland, A., & Brown, J. (2014). Coercive control in same-sex intimate partner violence. *Journal of Family Violence, 29*(1), 15-22.
- Gerbrandij, J., Rosenfeld, B., Nijdam-Jones, A., & Galietta, M. (2018). Evaluating risk assessment instruments for intimate partner stalking and intimate partner violence. *Journal of Threat Assessment and Management, 5*(2), 103-118.
- Graham, L. M., Sahay, K. M., Rizo, C. F., Messing, J. T., & Macy, R. J. (2021). The Validity and Reliability of Available Intimate Partner Homicide and Reassault Risk Assessment Tools: A Systematic Review. *Trauma, Violence, & Abuse, 22*(1), 18-40. <https://doi.org/10.1177/1524838018821952>
- Hall, M. (2021). Changes in New South Wales criminal justice since 2010. *Current Issues in Criminal Justice, 33*(2), 264-276.
- Hamberger, L. K., Larsen, S. E., & Lehrner, A. (2017). Coercive control in intimate partner violence. *Aggression and Violent Behavior, 37*, 1-11.

- Hardesty, J. L., Crossman, K. A., Haselschwerdt, M. L., Raffaelli, M., Ogolsky, B. G., & Johnson, M. P. (2015). Toward a standard approach to operationalizing coercive control and classifying violence types. *Journal of Marriage and Family*, 77(4), 833-843.
- Harris, B. A. (2020). Visualising violence? Capturing and critiquing body-worn video camera evidence of domestic and family violence. *Current Issues in Criminal Justice*, 32(4), 382-402.
- Hassouneh, D., & Glass, N. (2008). The influence of gender role stereotyping on women's experiences of female same-sex intimate partner violence. *Violence Against Women*, 14(3), 310-325.
- Jaffe, P. G., Johnston, J. R., Crooks, C. V., & Bala, N. (2008). Custody disputes involving allegations of domestic violence: Toward a differentiated approach to parenting plans. *Family Court Review*, 46(3), 500-522.
- Johnson, M. P. (1995). Patriarchal terrorism and common couple violence: Two forms of violence against women. *Journal of Marriage & Family*, 57(2), 283-294.
- Johnson, M. P. (2008). *A typology of domestic violence: Intimate terrorism, violent resistance, and situational couple violence*. Northeastern University Press.
- Johnson, M. P., & Leone, J. M. (2005). The differential effects of intimate terrorism and situational couple violence: Findings from the National Violence Against Women Survey. *Journal of family issues*, 26(3), 322-349.
- Joint Select Committee on Coercive Control. (2021). *Coercive control in domestic relationship*, Report 1/57, Retrieved on 18th January 2023 from <https://www.parliament.nsw.gov.au/ladocs/inquiries/2626/Report%20-%20coercive%20control%20in%20domestic%20relationships.pdf>
- Karystianis, G., Adily, A., Schofield, P. W., Greenberg, D., Jorm, L., Nenadic, G., & Butler, T. (2019). Automated analysis of domestic violence police reports to explore abuse types and victim injuries: Text mining study. *Journal of Medical Internet Research*, 21(3), e13067.
- Kuennen, T. L. (2007). Analyzing the impact of coercion on domestic violence victims: How much is too much. *Berkeley Journal of Gender, Law & Justice*, 22(2).
- Kuhn, M. (2014). *Futility analysis in the cross-validation of machine learning models*. *arXiv preprint arXiv:1405.6974*.
- Leung, F. & Trimboli, L. (2022) *Improving police risk assessment of intimate partner violence*. (Crime and Justice Bulletin No. 244).
- Michalski, J. H. (2005). Explaining intimate partner violence: The sociological limitations of victimization studies. *Sociological forum* (Vol. 20, No. 4, pp. 613-640).
- Mullainathan, S., & Spiess, J. (2017). Machine learning: an applied econometric approach. *Journal of Economic Perspectives*, 31(2), 87-106.
- Myhill, A. (2015). Measuring coercive control: What can we learn from national population surveys?. *Violence against women*, 21(3), 355-375.
- Myhill, A., & Hohl, K. (2019). The "golden thread": Coercive control and risk assessment for domestic violence. *Journal of interpersonal violence*, 34(21-22), 4477-4497.
- Otter, C. & Bosanko, M. (2022) *What is coercive control?* Retrieved on 18th January 2023 from <https://www.parliament.vic.gov.au/publications/research-papers/download/36-research-papers/14048-what-is-coercive-control>
- Pfitzner, N., Walkate, S., & McCulloch, J. (2022). Body-worn cameras: An effective or cosmetic policing response to domestic and family violence? *Criminology & Criminal Justice*, <https://doi.org/10.1177/17488958221108478>.

- Sheehy, E. (2018). Expert evidence on coercive control in support of self-defence: The trial of Teresa Craig. *Criminology & Criminal Justice*, 18(1), 100-114.
- Sherman, L. W. (2018). Policing domestic violence 1967–2017. *Criminology & Public Policy*, 17(2), 453-465.
- Straus, M. A. (1979). Measuring Intrafamily Conflict and Violence: The Conflict Tactics (CT) Scales. *Journal of Marriage and the Family*, 41(1), 75-88.
- Straus, M. A. (2011). Gender symmetry and mutuality in perpetration of clinical-level partner violence: Empirical evidence and implications for prevention and treatment. *Aggression and Violent Behavior*, 16(4), 279-288.
- Stark, E. (2007). *Coercive control*. Oxford University Press.
- Svalin, K., & Levander, S. (2020). The predictive validity of intimate partner violence risk assessments conducted by practitioners in different settings—a review of the literature. *Journal of Police and Criminal Psychology*, 35(2), 115-130.
- Tanha, M., Beck, C. J., Figueredo, A. J., & Raghavan, C. (2010). Sex differences in intimate partner violence and the use of coercive control as a motivational factor for intimate partner violence. *Journal of Interpersonal Violence*, 25(10), 1836-1854.
- Walby, S., & Myhill, A. (2001). New survey methodologies in researching violence against women. *British Journal of Criminology*, 41(3), 502-522.
- Walby, S., & Towers, J. (2018). Untangling the concept of coercive control: Theorizing domestic violent crime. *Criminology & Criminal Justice*, 18(1), 7-28.
- Walkate, S., & Fitz-Gibbon, K. (2021). Why criminalise coercive control? The complicity of the criminal law in punishing women through furthering the power of the state. *International Journal for Crime, Justice and Social Democracy*, 10(4), 1-12.
- Walkate, S., Fitz-Gibbon, K., & McCulloch, J. (2018). Is more law the answer? Seeking justice for victims of intimate partner violence through the reform of legal categories. *Criminology & Criminal Justice*, 18(1), 115-131.
- Watego, C., Macoun, A., Singh, D., & Strakosch, E. (2021). Carceral feminism and coercive control: When Indigenous women aren't seen as ideal victims, witnesses or women. *The Conversation*. Retrieved on 24th August 2022 from <https://theconversation.com/carceral-feminism-and-coercive-control-when-indigenous-women-arent-seen-as-ideal-victims-witnesses-or-women-161091>
- Whalley, E., & Hackett, C. (2017). Carceral feminisms: The abolitionist project and undoing dominant feminisms. *Contemporary Justice Review*, 20(4), 456-473.
- Wiener, C. (2017). Seeing what is 'invisible in plain sight': Policing coercive control. *The Howard Journal of Crime and Justice*, 56(4), 500-515.