

The dynamics of domestic violence: learning about the match

Dynamics of domestic violence

Dan Anderberg^{1,3,4*}, Noemi Mantovan², and Robert Sauer^{1,5}

Abstract: We present a dynamic lifecycle model of women's choices with respect to partnership status, labour supply and fertility when they cannot directly observe whether a given male partner is of a violent type or not. The model is estimated by the method of simulated moments using longitudinal data from the Avon Longitudinal Study of Parents and Children. The results indicate that uncertainty about a partner's abusive type creates incentives for women to delay fertility, reduce fertility overall, divorce more often and increase labour supply. We also study the impact of higher female wages, income support to single mothers, and subsidised childcare when the mother is working. While higher wages reduce women's overall exposure to abuse, both income support and subsidised childcare largely fail to do so by encouraging early fertility. The latter two policies instead increase the incidence of abuse towards mothers and hence the abuse exposure for children.

JEL Classification: J12, J13

Keywords: domestic violence, learning, fertility, ALSPAC

We thank seminar participants at Collegio Carlo Alberto, Stony Brook, IDC, CESifo, ISER, Oslo, Leicester, Trondheim, Cambridge, Kings College, Mannheim, Royal Holloway, Bangor, Trinity College Dublin, Toulouse, Ariel University, Ecole Polytechnique, IWCEE, ESPE, EEA, and SEHO for helpful discussions and comments. We are especially grateful to Michael Keane and Ahu Gemici for comments that greatly improved the paper.

1 Introduction

Freedom from violence is a fundamental human right. Yet violence by men towards their female partners is prevalent in every part of the world: reporting on the global prevalence of intimate partner violence WHO (2021) note that, among all ever-married/partnered women aged 15-49 world-wide, 27 percent have been subjected to physical and/or sexual violence from a current or former husband or male intimate partner at least once in their lifetime (24 percent for the UK), with far-reaching consequences for health, productivity, and well-being. Apart from its ubiquitous nature, domestic violence stands out as being the crime-category with the highest degree of repeat victimisation. For instance, in the UK – which is the focus of the current paper – while seven percent of all women aged 16-59 experienced domestic abuse in 2009/10, repeat victimisation accounted for more than three-quarters of all incidents of domestic violence (Flatley *et al.*, 2010).

Economics has recently seen a surge in research on domestic violence which has provided a wealth of useful insights. This research has focused on a range of environmental determinants of domestic abuse, including labour market conditions (Aizer, 2010; Tertilt and van den Berg, 2015; Anderberg *et al.*, 2016; Tur-Prats, 2021; Bhalotra *et al.*, 2021), educational attainment (Erten and Keskin, 2018), culture and social norms (Alesina *et al.*, 2021; Tur-Prats, 2019; González and Rodríguez-Planas, 2021; Guarnieri and Rainer, 2021), health and health innovations (Papageorge *et al.*, 2021), gender ratios (Amaral and Bhalotra, 2017), divorce laws (Stevenson and Wolfers, 2006; Garcia-Ramos, 2021), and, most recently, the COVID-19 pandemic (Piquero *et al.*, 2021; Berniell and Facchini, 2021).

The literature has further focused on understanding motives for and triggers of abuse, including instant gratification (Tauchen *et al.*, 1991), emotional cues (Card and Dahl, 2011), and instrumental abuse to change the victim's behaviour (Anderberg and Rainer, 2013) or to extract resources from the victim's family (Bloch and Rao, 2002; Haushofer *et al.*, 2019). Finally, there has been a number of studies of the effect of policy on the incidence of domestic abuse, including law enforcement policy (Iyengar, 2009; Aizer and Dal Bó, 2009), and welfare and cash-transfers policy (Angelucci, 2008; Bobonis *et al.*, 2013; Hidrobo and Fernald, 2013; Ramos, 2016; Hsu, 2017).

However, even with this flurry of contributions, a number of core questions – particularly of a dynamic nature – remain open. For instance, a question that has long been debated in the sociology and psychology literature is the dynamic link between a woman's labour supply and her exposure to abuse (Macmillan and Gartner, 1999; Riger and Staggs, 2004). This research has struggled with the fact that causality may go in both directions, and has been hampered by the use of relatively small and selective samples. Similarly, while there has been research into the relationship between domestic abuse and fertility, most of this research has focused particularly on abuse risk during pregnancy (Jasinski, 2004; Bowen *et al.*, 2005). Finally, perhaps the most obvious dynamic response to abuse is whether or not a woman leaves her partner (Enander and Holmberg, 2008; Bowlus and Seitz, 2006).

The aim of this paper is to construct and estimate a dynamic lifecycle model of women’s choices with respect to partnership status, fertility and labour supply in an environment where they are at risk of abuse from their partners and to use the estimated model to predict responses to changes in the economic environment, including policy. Our primary focus will be on modelling women’s behavior in order to better understand how they navigate such an environment. As we will argue below, the data strongly suggest that there is heterogeneity among men. A starting point for our model is that women generally will, when meeting a new potential partner, not be able to directly observe his complete nature. Instead they will infer this by observing his behaviour over time. In our model we incorporate such learning in the simplest possible Bayesian form. A man either has a “violent nature” or a “non-violent nature” where the former type is abusive with a high frequency and the latter only rarely. A woman holds beliefs over her partner’s type which she updates based on observing his behaviour. If she experiences abuse her expectations about what the future within the relationship would hold worsen, potentially making her choose divorce, and a change in labour supply, and/or fertility. In addition, a woman may delay fertility within a relationship until she is reasonably certain about her partner’s nature.

Our modelling of male behaviour – specifically that of violent men – is as flexible as possible in order to be consistent with alternative theories of abusive behaviour, for instance exposure, identity, and bargaining theory. In particular, we model the abusive behaviour of violent males as following a behavioural rule that we include as part of the estimation. We specify that several factors, including e.g. her labour supply, contribute to “tension” which, when combined with random cues can lead to abuse “triggers”, some of which can be averted through self-control. In this way our modelling of male behaviour is agnostic, for instance allowing abuse to partially reflect strategic behaviour without necessarily assuming it.

To estimate the model, we use data from the Avon Longitudinal Study of Parents and Children (ALSPAC), a local longitudinal study that has followed a set of children – and their parents – from birth. Our sample population will include close to 4,000 ALSPAC first-time mothers who are followed for seven years starting from the study pregnancy.¹ Importantly, the survey contains annual measures of intimate partner abuse, and we observe partnership, labour supply and fertility choices. The ALSPAC data, by surveying pregnant women, involves choice-based sampling, which as argued by Manski and Lerman (1977) and Cosslett (1981), has both advantages and drawbacks: it confers efficiency gains when some alternatives of particular interest are otherwise infrequently chosen but also poses challenges in terms of finding a suitable estimator that accounts for the non-random sampling. In our case child-bearing is an infrequent choice of central interest informing about women’s investments within their relationships. In this respect, ALSPAC provides a clear advantage over other potential data sets of more

¹Our focus on first-time mothers is to capture women at the early stages of their relationships where more learning can be expected to take place.

general female populations.² We handle the non-random sampling by using the method of simulated moments, simulating lifecycle paths for a population of women and then computing moments on the simulated data selected using the ALSPAC sampling approach.

We model a woman's choice of partnership status, labour supply, and child-bearing from her late teens until the end of her fertile period. As such, our model builds on an established literature developing lifecycle models of family decisions.³ The relationship between our work and two contributions to this literature are worth noting in more detail. Brien *et al.* (2006) focus on the choice between marriage and cohabitation and a couple jointly learns about the true match quality of their match. Our learning setting is on the one hand simpler: women learn their partner's type with only a binary type space, and update beliefs based on abuse which is observed in the data. On the other hand, by endogenising fertility and labour supply we study key behavioural responses to learning beyond partnership decisions. Bowlus and Seitz (2006) is the only contribution to date to estimate a lifecycle model with domestic violence. In their model, men rationally decide on abuse based on their preferences for violence. However, as women always know their partner's abuse preferences there is no learning. Moreover, fertility is treated as exogenous.

Our results indicate that violent men are high-frequency repeat abusers. As a result learning is quite fast – within a few years most women will be quite certain about the nature of their partner – creating a strong incentive to delay fertility within a new relationship by a few years. The uncertainty and learning also mean that fertility is lower overall and divorces are higher than would be the case had male types been immediately observable. Our results further indicate that a woman is less at risk of abuse from a violent partner when she participates in the labour market. In contrast we find that an increase in a woman's earnings capacity has only a modest effect of reducing the rate of abuse by a violent partner.

The latter finding does not imply that higher female wages would not reduce the equilibrium incidence of abuse. Indeed, our counterfactual simulations highlight a clear effect, but this effect occurs mainly through key behavioural responses: divorcing more frequently, delaying fertility, and working more frequently. In other words, better labour market opportunities imply that women are less likely to become trapped in abusive relationships as they are less likely to have children early in relationships and as they are financially better placed to leave bad relationships.

We also explore the potential effects of (i) an increase in the income support available to single mothers, and (ii) subsidised childcare available to households where the mother is working. In each of these two scenarios fertility is encouraged – with the former particularly encouraging single motherhood and the latter particularly encouraging fertility early in relationships. In addition, income support for single mothers leads to lower labour supply over the lifecycle. As a result, more generous income support

²It should also be noted that longitudinal data on domestic violence are extremely rare and most other datasets are small and highly selective in other ways, for instance often focusing on women seeking support after experiencing abuse. A drawback in our case is that the data will not have any information about the behaviour of women who choose to remain childless.

³Key contributions include van der Klaauw (1996), Francesconi (2002), Keane and Wolpin (2010), and Gemici and Laufer (2014).

provided to single mothers perhaps somewhat surprisingly fails to reduce abuse exposure overall. In contrast, subsidised childcare encourages labour force participation which mitigates the effect of early child bearing on exposure to abuse, leading to only a minor overall effect on abuse. A worrying unintended consequence of both policies however is that they lead to a higher incidence of abuse experienced by mothers in particular, implying that children are more likely to be exposed to abuse between their parents. These results highlight the value of having a full dynamic model that consider not only the direct effects, but also the overall effects via key dynamic choices.

The paper is outlined as follows. Section 2 describes the ALSPAC sample and some of the key empirical facts that our empirical model will seek to replicate. Section 3 describes the model, starting with a simple illustrative version before outlining the full empirical version. Section 4 outlines the estimation approach while Section 5 reports the model fit and the parameter estimates. Section 6 presents the counterfactual experiments with perfect information, elimination of the gender pay gap, more generous child support and subsidised childcare. Section 7 concludes.

2 The ALSPAC Data and Key Empirical Features

The Avon Longitudinal Study of Parents and Children (ALSPAC), also known as “Children of the 90s” is a local UK cohort study conducted in the former England county of Avon. The initial recruits were pregnant women with estimated dates of delivery between April 1991 and December 1992.⁴ While first and foremost a child development survey, ALSPAC also repeatedly surveyed the mothers of the study children (and their partners). Our female sample population is hence drawn from the set of mothers of the ALSPAC children and we exploit the fact that the mothers were surveyed roughly annually about key events in their lives, including their experience with abuse. These questions were included in the mothers-surveys up until when the survey child was about 6 years old, yielding a maximum of seven observation years for each female respondent.⁵

As we are particularly interested in modelling partnership dynamics in which women form and update beliefs about the nature of the partners we want to focus on female respondents with generally short partner durationships at the start of the study. For this reason we restrict our sample to first-time mothers. An advantage to using longitudinal data around child-birth is that this is a key period when women’s decisions regarding further fertility and if and when to return to work are particularly salient.

The ALSPAC data provides a unique opportunity for studying the dynamics of domestic abuse, both in terms of incidence and behaviour. To this end, we provide an extensive description of the sample used and the key dynamic patterns in the data in online Appendix A. Here we provide a brief overview of key

⁴Ethical approval for the study was obtained from the ALSPAC Ethics and Law Committee and the Local Research Ethics Committees. For a detailed description of the ALSPAC cohort, see Boyd *et al.* (2013).

⁵The survey mothers completed multiple questionnaires during their pregnancy, one of which included the key questions on partner abuse. Post-birth they were asked to complete surveys with the abuse questions when the study child was aged 8, 21, 33, 47, 61 and 73 months respectively. After that the key abuse-related questions were no longer regularly asked.

features. ALSPAC recruited 14,541 pregnant women who returned at least one questionnaire or attended at least one clinic. After restricting the sample (see Appendix A for details) to first-time mothers of white ethnic origin, with complete information on basic demographics characteristics and at least one completed post-pregnancy questionnaire, we are left with 3,960 women, with a total of 22,204 person-year observations.

2.1 Marriage, Births and Labour Supply

In our analysis we make no distinction between formal marriage and cohabitation as effective learning about the nature of a partner can be expected to start when first living together.⁶ Hence we will use being “married” as synonymous to “living with a partner”. The vast majority – 96 percent (Table A1) – of the women in the sample live with a partner (are “married”) at baseline, that is during the ALSPAC pregnancy. However, we observe both separations and entries into new partnerships (a 1.8 percent annual separation rate among with partners and a 11.6 percent annual partnership entry rate among those single, Table A2).

The restriction to first time mothers implies that most of the women in the sample are in relatively “new” relationships at baseline: close to 60 percent of women with partners at baseline have lived with those partners for no more than 3 years and the vast majority of them were in their 20s at baseline (Figure A1).

All women, per construction of the data, have a birth between the first and the second period in the data. For about 70 percent of the women in the sample we also observe the birth of at least one further child. The observed annual fertility rate after the second period is 19.5 percent (Table A2). The vast majority of women are not working immediately following the birth of the ALSPAC child. However, over time many return to work. Across all periods we have a 22.2 percent rate of transition from not working to working, either part- of full-time (Table A2). We delineate three levels of educational attainment – “low”, “medium” and “high” – of roughly equal-size based on a standard mapping of academic qualifications into National Vocational Qualification (NVQs) equivalents used by the Office for National Statistics (Table A1). We impute hourly wages based on information on current or most recent occupation, and we show that wages increase both with age and with qualification level, for both women and men (Table A3).⁷

2.2 The Incidence of Abuse

A key issue is the measurement of abuse. The literature typically advocates strict objective measures (Aizer, 2010; Tertilt and van den Berg, 2015) which is natural in contexts where the research aim is to

⁶Learning could of course start during “dating”. Data on the duration of the pre-coresidential phase of relationships is scarce in social surveys. Evidence from both the US and Germany (Yau, 2019, Schnor, 2015), including on individuals born through the 1960 and 1970a, suggest that couples who eventually moved in together dated for on average around 1.5 years.

⁷Average hourly wages by gender and occupation was obtained as a tabulation from the Labour Force Survey 1993-1999 (ONS, 1999).

understand the effect of various factors on the incidence of abuse. Our aim, in contrast, is to understand a woman’s behavioural responses to her experience of abuse and the associated changes in her beliefs about the nature of her current partner. In line with this aim, we make use of a self-reported measure of physical and emotional abuse: whether the respondent reports that the partner has been “physically or emotionally cruel” to her since the last survey. Overall, we find that, across all observations, 2.0 and 7.3 percent of women report experiencing physical and emotional abuse respectively over the past year (Table A4). For our analysis we combine the two into a single indicator of abuse of “any kind” with an average incidence rate of 7.8 percent. While the abuse-related questions are less specific than ones used in many dedicated domestic violence survey modules, it can be shown that the estimated incidence of physical and emotional abuse in our sample is very similar to the best available evidence from the Crime Survey for England and Wales. We also show that the measured abuse is, as expected, higher among younger women than among older women, and higher among the less qualified women than among the more qualified women (Figure A2).

The value of using longitudinal data becomes clear when we consider the persistence of abuse. We find evidence of a very high persistence: of women reporting experiencing abuse in some period t , 45.7 percent report experiencing abuse again in the following period (Table A4). A key modelling challenge is capturing this high rate of persistence of abuse experienced by some women.

The incidence of abuse decreases monotonically with partnership duration, from 13 percent in relationships with 0-1 years duration down to 7 percent in relationships with durations of 6 or more years (Figure A2). This relationship is of course highly endogenous as women select out of relationships based on their experience. As we will show below, this duration-pattern will be instrumental both to understanding the prevalence of violent males and to women’s responses to abuse.

As noted above, key discussion in the literature is the relationship between labour supply and abuse. A complication is that this relationship may be bi-directional. Using the longitudinal nature of the data we highlight that there is a U-shaped relationship between labour supply at $t - 1$ and abuse experienced between $t - 1$ and t : 8.4 percent of women who were not working a year ago report experiencing abuse over the past 12 months, while those who were working part-time and full-time report abuse at rates of 7.2 and 9.3 percent respectively (Figure A2).

2.3 Choices Following Abuse

The longitudinal nature of the data also allows us to explore how women’s observed choices with respect to partnership status, child-bearing and labour supply relate to their experience of abuse.

A controversial question is whether women who experience abuse stay in their relationships. We find that the rate of separation following a period of abuse is about five times higher than the corresponding rate following a period without abuse (Table A5). In other words, the data suggests that women do

systematically leave abusive partners at a substantially higher rate. However, at the same time it also suggests that women who experience abuse are, in any given period, far more likely to stay than to leave.

What is unique about the current data is that it allows us to consider the dynamic relationship between abuse and fertility. Here we find that the fertility hazard is reduced by about a third after a period of experiencing abuse. This suggests that a strong response to abuse is to reduce fertility (Table A5). In general, the dynamic relationship between fertility and abuse will be a central feature of the estimated full model.

When we look at how women may respond to abuse in term of their labour supply – relating the experience of abuse between $t - 1$ and t to labour supply status at time t – the data suggests an increase in labour force participation though the response appears rather modest (Table A5).

3 Model

We develop a model of the behaviour of women in an environment where there is heterogeneity among males with respect to their propensity to engage in abuse: some men have a “violent” nature and some do not. We model abuse as non-strategic occurrences of loss of control, though possibly influenced by an underlying gender-tension. A woman who meets a new prospective partner does not directly observe his type; instead she forms a belief which she updates based on her observations of his behaviour.

Women also choose labour supply and fertility. The interaction between learning and fertility is particularly interesting as it leads to the possibility that a woman becomes “trapped” in abusive relationships. Once a woman has children either childcare costs have to be incurred or she will have to lower her labour supply and forego earnings. This makes it more financially difficult for her to divorce once children are present and as a consequence she will be more prone to stay even if that means suffering abuse. To avoid this she can delay fertility until she knows her husband’s type better and also use that delay to gain further labour market experience. Thus delaying fertility and building labour market experience act as a type of insurance that increases her flexibility in case she discovers that her partner is of the violent type.

Before presenting the full empirical model we will begin by presenting a simple illustrative version that ignores labour supply and fertility but introduces the core learning structure. In particular, we will use this simple model to highlight how the main type- and learning structure allows us to replicate key features in the data relating to the incidence of abuse.

3.1 A Simple Illustrative Version

Consider a population of women who are facing an infinite time horizon, $t = 1, 2, \dots$, and who in any given period t are either single or married, $m_t \in \{0, 1\}$. Normalise the instantaneous utility of being

single to zero and the systematic utility of being married, denoted ψ^m , to unity. Additionally, let ε_t^m be a temporary marriage utility shock which is normally distributed with zero mean and variance σ_m^2 .

A currently married woman can either remain married or divorce. Single women receive marriage offers at rate ζ from randomly drawn prospective partners. Men are of two possible types, $r \in \{0, 1\}$, who differ in their propensity for abuse: a “non-violent type” ($r = 1$) and a “violent type” ($r = 0$). A woman receiving a marriage offer does not observe the male’s type. Let $\phi_b = E[r] \in (0, 1)$ be the probability of him being non-violent. ϕ_b thus also represents the woman’s initial beliefs.

Let $z_t \in \{0, 1\}$ indicate abuse in period t and let $\chi_r = \Pr(z_t = 1|r)$ denote the probability that a type- r male is abusive. We assume that $0 < \chi_1 < \chi_0 < 1$. A woman updates her beliefs based on her husband’s behaviour. Under standard Bayesian updating, a woman who holds beliefs ϕ_{t-1} going into period $t - 1$ and who *does not* experience any abuse in that period will hold the next period belief

$$\phi_{t|z_{t-1}=0} = \frac{\phi_{t-1}(1 - \chi_1)}{\phi_{t-1}(1 - \chi_1) + (1 - \phi_{t-1})(1 - \chi_0)}, \quad (1)$$

whereas if she *does* experience abuse her next period belief will be

$$\phi_{t|z_{t-1}=1} = \frac{\phi_{t-1}\chi_1}{\phi_{t-1}\chi_1 + (1 - \phi_{t-1})\chi_0}. \quad (2)$$

Abuse is associated with a instantaneous disutility $\psi^z > 0$. Hence the expected disutility from abuse in period t for a married woman with current beliefs ϕ_t are $\pi(\phi_t)\psi^z$, where $\pi(\phi_t) = \phi_t\chi_1 + (1 - \phi_t)\chi_0$. Letting δ denote the discount rate, the model can be solved using basic dynamic programming. In particular, there will be a present discounted value $V^m(\phi_t)$ associated with entering a period as married with belief ϕ_t and a value V^s associated with entering a period as single.⁸ A woman with current belief ϕ_t will choose marriage over singlehood if

$$\psi^m + \varepsilon_t^m - \pi(\phi_t)\psi^z + \delta \left[\pi(\phi_t) V^m(\phi_{t+1|z_t=1}) + (1 - \pi(\phi_t)) V^m(\phi_{t+1|z_t=0}) \right] \geq \delta V^s, \quad (3)$$

which means that there will be a threshold ε_t^m below which she will divorce. Moreover, this threshold value will be a decreasing function of ϕ_t , creating a positive link between abuse and divorce.

The model can be easily calibrated and doing so is instructive as it illustrates how some of the core parameters are identified from key empirical moments in the data. Such a calibration is shown in Table 1.⁹ The observed overall annual divorce rate is 1.8 percent (Table A2), but it is substantially higher following abuse than following non-abuse (7.6 percent and 1.4 respectively – see Table 2 below). These moments are intuitively informative about ψ^z and σ_m^2 . The fact that women who experience abuse are

⁸Formally, $V^m(\phi_t)$ and V^s satisfy

$V^m(\phi_t) = E\varepsilon_t^m \left[\max \left\{ \psi^m + \varepsilon_t^m - \pi(\phi_t)\psi^z + \delta \left[\pi(\phi_t) V^m(\phi_{t+1|z_t=1}) + (1 - \pi(\phi_t)) V^m(\phi_{t+1|z_t=0}) \right], \delta V^s \right\} \right]$ and $V^s = \zeta V^m(\phi_b) + \delta(1 - \zeta)V^s$ respectively.

⁹In this calibration we have set the discount parameter to $\delta = 0.95$.

markedly more likely to divorce than those who don't, while at the same time also being more likely to remain married than to divorce, indicate that the disutility term ψ^z is substantial but also below the (normalised) systematic utility of marriage, ψ^m . The frequency of divorce after non-abuse is informative about the size of the match quality shocks.

Table 1: A calibration of the simple illustrative model.

Calibrated Moments		Parameter	
Singlehood duration (yrs) (m_1)	6.94	$\zeta = 0.140$	Meeting rate
Divorce rate after non-abuse (m_2)	0.014	$\sigma_m^2 = 2.615$	Match quality shock variance
Divorce rate after abuse (m_3)	0.076	$\psi^z = 0.280$	Abuse disutility
Abuse rate overall (m_4)	0.072	$\phi_b = 0.752$	Proportion non-violent
Abuse persistence (m_5)	0.457	$\chi_1 = 0.023$	Abuse rate: non-violent males
Abuse rate at zero duration (m_6)	0.175	$\chi_0 = 0.695$	Abuse rate: violent males

Notes: The left part of the table presents empirical moments. All abuse-related moments come from the ALSPAC data and will be presented in further detail in Table 2 below. For duration of singlehood, see text. The discount rate δ is set of 0.95. The right part presents the parameters that make the steady state of the simple illustrative model match the empirical moments.

The duration of singlehood relates to ζ .¹⁰ We do not have any direct measure of this duration in the data, but based on the rate of entry into new partnerships (Table A2) and on the average age at first birth (see Table B1), we can approximate this as about seven years. Not all marriage proposals in this simple model are accepted, but most are, whereby ζ is closely above the inverse of this duration.¹¹

The final three parameters – χ_0 , χ_1 and ϕ_b – can be set to match key empirical moments of abuse. First, in order to generate a high abuse persistence (Table A4), it has to be that some males are high repeat abusers, in particular χ_0 has to be well above 0.5.¹² Second, the incidence at the start of a relationship (Figure A2, panel c) is a direct linear combination of three parameters, $\phi_b\chi_1 + (1 - \phi_b)\chi_0$, whereas the overall incidence rate (Table A4) is significantly lower due to selective divorces.

The calibrated parameters are highlighted in Table 1. As we will see, the insights from this simple exercise will carry over to the main model below. For instance, the estimated main model will suggest a similar size of disutility of abuse relative to the systematic utility of marriage, and similar values for the abuse rates by male types, and a similar type frequency among males.

As a further validity check, we can follow Andrews *et al.* (2017) and measure the sensitivity of the calibrated parameters with respect to the empirical moments. We provide further details of the approach


¹⁰Reassuringly, the duration of singlehood, and hence the hazard rate for entry into partnership, is very similar to, for instance, Coles and Francesconi (2019).

¹¹Note that a woman is as likely to accept a new partner as she is to remain married at belief ϕ_b .

¹²Our model explains persistence of abuse by permanent heterogeneity among men, not by direct state-dependence in abusive behaviour. As a check on this, we verify that the persistence of abuse between t and $t - 1$ is similar to that between t and $t - 2$. See Table A5 (Panel D). We also explore the validity of assuming that women draw random men by checking that, among those who divorce at t , those who do so with an experience of past abuse are neither more likely to report abuse in a future relationship nor more or less likely to enter any new relationship.

in Appendix C and show how it can generate an approximation of elasticity of each of the six estimated parameter with respect to each of the empirical moments.

Figure 1: Sensitivity of calibrated parameters to empirical moments.



Notes: The figure shows the sensitivity, in elasticity form, of the calibrated parameters with the six empirical moments. The parameters and moments ($m_1 - m_6$) are described in Table 1. For further details of the method for approximating the elasticities, see Appendix C.

The resulting elasticities are shown in Figure 1. This confirms that the estimate of each parameter is responsive to some empirical moment. For instance, the estimated meeting rate ζ is mainly and negatively related to the singlehood duration; the estimated match quality variance σ_m^2 is positively related to the no-abuse divorce rate and the singlehood duration; the estimated abuse disutility ψ^z is strongly related to the post-abuse divorce rate; the estimated proportion non-violent males ϕ_b relates negatively to the overall abuse rate; finally the estimated probabilities of abuse from the two types of men, χ_1 and χ_0 , respond strongly to the abuse persistence and abuse risk at the beginning of a relationship. Figure 1 suggests the parameter estimates are somewhat sensitive to the final two moments in particular.

The current simple model can be used to gauge the speed of learning. Given the sharply different behaviours of violent and non-violent males, learning will generally be quite fast. Consider for instance a large set of women who all marry randomly selected partners at some time t and remain married for at least three periods. Using the calibrated values of χ_0 , χ_1 and ϕ_b , after three periods, about two-thirds of the women will have not experienced any abuse and will hold a belief that the partner is non-violent that is well-above 0.95. Conversely, about 20 percent of the women will have experienced two or three periods of abuse and will hold a belief that is below 0.01. In other words, within three periods of

marriage, the vast majority of women will be nearly certain about the true nature of their partners, in either direction.

3.2 The Full Empirical Model

While the illustrative model was useful for setting the stage for the learning environment and for highlighting how heterogeneity among men is essential for understanding the dynamic patterns of the incidence of abuse, it is also limited due to its focus on a stationary environment. In order to build in key dimensions, such as fertility and the accumulation of work experience, and in order to make our model more useful for policy analysis, we need a lifecycle model.

3.2.1 Setup

The full version that we take to the data models women's choices with respect to marital status, employment status and child-bearing from the time of entry into adulthood until the end of their fertile period, age 16 to 44, a total of $T = 29$ periods. In each period $t = 1, \dots, 29$ there are three mutually exclusive employment states $k_t \in \{0, 1, 2\}$, representing not-working, working part-time and working full-time respectively. As before $m_t \in \{0, 1\}$ indicates whether the woman lives with a male partner ("married") or not, and we let $f_t \in \{0, 1\}$ indicate the choice whether or not to conceive a child at time t .

Each woman maximises the present value of her lifetime utility, discounted at rate δ . The utility flow in period t is specified as

$$U_t = \frac{\mu^{k_t} C_t^{1-\lambda}}{1-\lambda} + (\Psi_t^m - \bar{\Psi}_t^z) m_t + \Psi_t^n, \quad (4)$$

where C_t is her level of consumption, μ^{k_t} varies with the employment state k_t , and λ is the parameter of relative risk aversion. μ^0 is normalised to unity while μ^1 and μ^2 are constrained to the unit interval to capture disutility of work effort. The following term, which is enjoyed by the woman only if she chooses to be married in period t , includes the direct utility of marriage Ψ_t^m and the expected disutility from abuse $\bar{\Psi}_t^z$. The final term, Ψ_t^n , captures the direct utility of children. The Ψ -terms will be further specified below.

The unit of time is taken to be a year, so consumption and earnings represent annual values. The consumption enjoyed by the woman at time t is

$$C_t = \begin{cases} \tau (w_t + w_t^h - c_t) & \text{if } m_t = 1 \\ w_t - c_t & \text{if } m_t = 0 \end{cases}, \quad (5)$$

where w_t and w_t^h are her own and her husband's annual earnings at t respectively, τ is an income sharing parameter, and, c_t represents annual child-related costs and incomes, specified further below.

3.2.2 Wages, Experience and Child-Related Costs

When not working the woman receives a fixed basic unearned income $w^0 > 0$. If she is in work, her earnings associated with part- and full-time work are

$$w_t^k = \exp\left(\beta_0^k + \beta_1^k a + \beta_2^k x_t + \beta_3^k x_t^2 + \varepsilon_t^k\right), \text{ for } k = 1, 2, \quad (6)$$

respectively, where $a \in \{0, 1\}$ is a fixed individual characteristic that captures permanent heterogeneity among women in earnings capacity and where x_t measures her accumulated work experience. A woman's permanent productivity type a is assumed to be stochastically related to her observed educational attainment, which, as described in Section 2, is either “low”, “medium”, or “high”, $q \in \{0, 1, 2\}$. We specify the relationship between q and a to be logistic,

$$\frac{\Pr(a = 1|q)}{\Pr(a = 0|q)} = \exp(\beta_0^a + \beta_1^a d_{q=1} + \beta_2^a d_{q=2}), \quad (7)$$

where d_q is a dummy for educational attainment level q and where low educational attainment is the base category.

Work experience, which is accumulated according to

$$x_{t+1} = x_t + k_t, \quad (8)$$

starts from the initial condition of zero. Her work experience thus increases by one unit if she works part-time and by two units if she works full-time. Finally, the part-time and full-time wage offers at time t include distinct temporary productivity shocks, ε_t^k , for $k = 1, 2$.

The husband's earnings in equation (5) is specified in a similar way as

$$w_t^h = \exp\left(\beta_0^h + \beta_1^h a + \beta_2^h t + \beta_3^h t^2 + \varepsilon_t^h\right), \quad (9)$$

where ε_t^h is also a temporary productivity shock. The presence of the woman's own permanent productivity type a in the husband's wage offer equation (9) captures a systematic spousal wage correlation thus representing marital sorting on ability. Since men are assumed to always be working (full-time) in our model, their work experience increases linearly with time t .

The distribution of the temporary productivity shocks is joint normal, $(\varepsilon_t^1, \varepsilon_t^2, \varepsilon_t^h) \sim N(\mathbf{0}, \mathbf{\Sigma})$ with covariance matrix $\mathbf{\Sigma} = \mathbf{A}\mathbf{A}'$ where \mathbf{A} is the Cholesky decomposition. \mathbf{A} is restricted for identification reasons so that

$$\mathbf{A} = \begin{bmatrix} a_{11} & 0 & 0 \\ a_{21} & a_{22} & 0 \\ a_{h1} & 0 & a_{hh} \end{bmatrix}. \quad (10)$$

The child-related costs and incomes c_t have two basic components. The first component is childcare costs. The maximum childcare costs are assumed to be quadratic in the number of children. A fraction ρ^{k_t} of the maximum childcare cost is incurred at labour supply level k_t . We normalise $\rho^2 = 1$ – that is, the full child-care cost is incurred when working full-time – whilst ρ^1 and ρ^0 are constrained to the unit interval and estimated. The second component of c_t is income support that accrues to single mothers. Such income may come from alternative sources, including benefits and child-support payments from the biological father.¹³ Given the potential multiple sources, we will model child-related income to single mothers in the simplest possible way as a quadratic function of the number of children and include it in the estimation. Hence we specify the two components of c_t as follows

$$c_t = \rho^{k_t} (\beta_1^{cc} n_t + \beta_2^{cc} n_t^2) - (\beta_1^{ci} n_t + \beta_2^{ci} n_t^2) (1 - m_t), \quad (11)$$

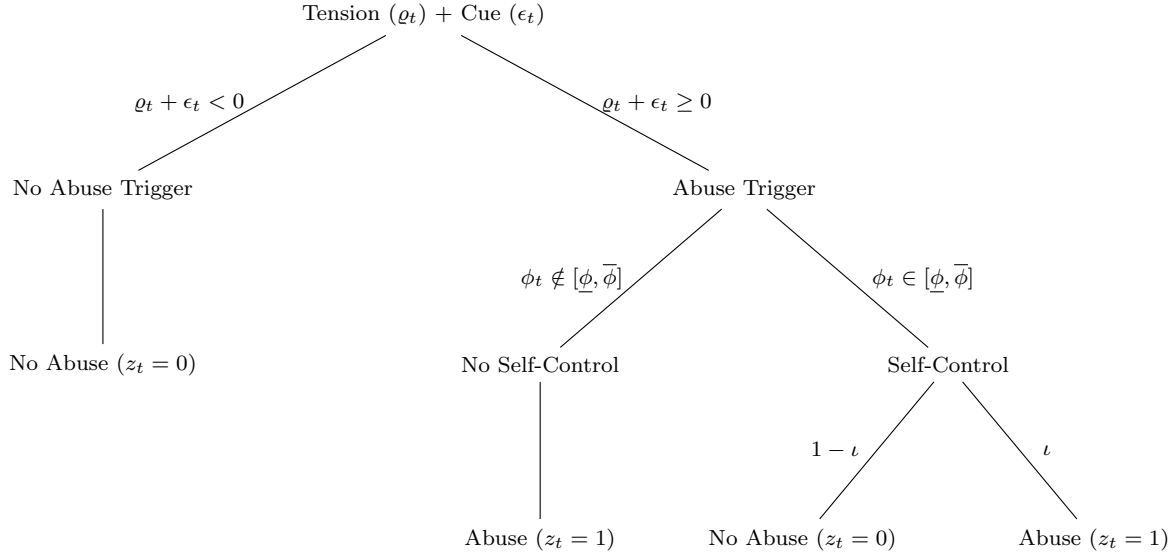
where the first term enters positively as it represents a cost and the second negatively as it represents income.

3.2.3 Abuse Risk

We continue to assume that there are two possible types of males, $r \in \{0, 1\}$. In the simple model version above, the probability of abuse only depended on the male’s type, being χ_0 for any *violent* male ($r = 0$) and χ_1 for any *non-violent* male ($r = 1$) with $\chi_0 > \chi_1$. In the full model we substantially enrich the specification of the probability of abuse by violent males, whilst we retain the assumption of a (small) fixed abuse probability for the non-violent males.

There are two key challenges in relation to modelling abuse behaviour. First, there are several alternative theories to explain such behaviour. Second, there is a general lack of consensus regarding the extent to which abuse reflects rational decision-making by perpetrators who have full agency over their actions. Our response to these challenges is to adopt a flexible specification that (i) allows us to capture the essential features of *alternative theories* of abuse, and (ii) allows for the possibility that violent males have *some agency* over their behaviour and follow a behavioural rule that would be broadly consistent with most plausible models of male utility-maximisation. To achieve (i) whilst retaining the core stochastic feature of abuse, we model “abuse triggers” as a combination of an underlying systematic “tension” and a randomly occurring “emotional cue”, where the former has multiple possible determinants. We will capture (ii) by allowing for the possibility that violent males can exert self-control that reduces the proportion of abuse triggers that feed through to realised abuse, and that they do so according to a simple behavioural rule that we motivate further below. Figure 2 illustrates our specification.

¹³During the period of study, “Income Support” (IS) was the main out-of-work benefit in the UK, with a maximum benefit that depended on the number and ages of children and that also included a lone-parent premium. Eligibility for IS was conditional on not working more than 16 hours/week. The in-work benefit system at the time was “Family Credit” (FC) which was designed for families with children where at least one person is working more than 16 hours/week. Lone mothers were a main recipient group for both IS and FC.

Figure 2: Specification of abuse risk from violent ($r = 0$) partners.

Consider first the specification of tension ϱ . We model ϱ so as to include key elements of exposure-, bargaining-, gender-identity and stress theory. Basic exposure theory in criminology (Dugan *et al.*, 2003) emphasises presence in high-risk places, which in the case of domestic abuse is inside the home; as a result exposure theory predicts a negative relationship between labour supply and abuse incidence: the more she works, the less time she spends in the household and the lower is exposure to abuse. Bargaining theory (Farmer and Tiefenthaler, 1997; Aizer, 2010) would suggest that improvements in her relative earnings capacity would strengthen her bargaining position and thereby reduce the incidence of abuse. Conversely, however, gender identity theory (Macmillan and Gartner, 1999) would suggest that the woman working more and/or having a higher earnings capacity – either in absolute or relative terms – may challenge the male’s identity and thereby increase tension. We further account for potential stressors to influence tension (Roberts *et al.*, 2011). Specifically, we allow tension to depend on the number of children and on time. A direct influence of time could for instance reflect an ability to cope with stress and frustrations that increases with age.¹⁴ Hence we will model current tension in a general fashion as depending on current labour supply, earnings capacities, number of children (linear and quadratic), and time,

$$\varrho(k_t, \bar{w}_t^2, \bar{w}_t^h, n_t, t) = \sum_{k=0,1,2} \chi_0^k d_{k_t=k} + \chi_0^w \log(\bar{w}_t^2) + \chi_0^h \log(\bar{w}_t^h) + \chi_0^n n_t + \chi_0^{n^2} n_t^2 + \chi_0^t t, \quad (12)$$

¹⁴Gottfredson and Hirschi (1990) provide a seminal discussion of the link between level of self-control, threshold for frustration, and tendency to resort to violence as a means to resolving conflict. In their study, they however argue that an individual’s capacity for self-control develops at a young age, by ages eight to ten, and then remains relatively stable over the life-course.

where $d_{k_t=k}$ is a labour supply dummy indicating whether $k_t = k$, and where \bar{w}_t^2 is the deterministic part of the woman's potential full-time earnings and \bar{w}_t^h correspondingly is the deterministic component of the husband's current earnings.¹⁵ Adding to the current systematic tension ϱ_t is a random emotional cue, denoted ϵ_t , which we take to be i.i.d. extreme value distributed across individuals and time periods. An *abuse trigger* is assumed to occur when $\varrho_t + \epsilon_t \geq 0$.

Modelling potential self-control by perpetrators of abuse is also challenging for several reasons. First, it is not clear how efficiently perpetrators are able to control their own behaviour. Second, in the same way that it is not clear to what extent perpetrator behaviour reflects rational, strategic and sophisticated behaviour. We will tackle these challenges by empirically modeling and estimating a simple specification in which control is imperfect and where self-control effort follows a simple behavioural rule that we argue would be broadly consistent with alternative assumptions about perpetrator decision-making.

First, we will assume that a perpetrator's self-control effort is a binary action and that self-control, when exerted, is generally imperfect. Specifically, we will assume that if a violent male *does not* exert self-control effort, then every "abuse trigger" leads to the occurrence of abuse.¹⁶ In contrast, if a violent male *does* exert self-control, then only some proportion $\iota \in [0, 1]$ of "abuse triggers" lead to the occurrence of abuse.

Second, we want to consider *when* a utility-maximising violent male would choose to exert self-control, assuming that doing so is privately costly. Given that self-control is assumed to be a binary action, any model of the male's behaviour would imply that they, in equilibrium, exert self-control within some subset of state-space. As we are not interested *per se* in estimating the males' preferences and effort costs, but merely accounting for their possible self-control behaviour, this offers an attractive modelling alternative: to directly estimate in which states a violent male exerts self-control effort. This approach thus encompass wide range of possible assumptions about violent males' decision-making – e.g. whether they are myopic or forward-looking, and whether or not they exhibit full strategic sophistication.

Indeed, we will simplify this further by focusing on one key dimension of the state-space, namely the wife's beliefs. This focus on her beliefs follows naturally if we assume that he wants to remain married and associates a worsening of the wife's beliefs with an increased risk of divorce (as we saw in the simple model). We would then expect that a violent male will exert self-control when the abuse realisation would have a relatively large impact on the wife's beliefs. This will generally be the case when her current beliefs ϕ_t fall within some intermediate range $[\underline{\phi}, \bar{\phi}]$ where $0 < \underline{\phi} < \bar{\phi} < 1$.¹⁷ In contrast, when she either firmly believes that he is of the violent type (ϕ_t is close to zero) or that he is of the non-violent type (ϕ_t is close

¹⁵The deterministic components \bar{w}_t^2 and \bar{w}_t^h are thus (6) (for $k = 2$) and (9) evaluated with ε_t^2 and ε_t^h each set to zero. These thus capture current earnings capacity as arising from ability and work experience.

¹⁶Note that this is without loss of generality: if, in the absence of self-control, some proportion of triggers did not lead to abuse, this would be equivalent to there being a proportionately lower trigger frequency.

¹⁷To see this more formally, consider how the wife's beliefs change after experiencing a period of abuse. Using (2) this change takes the form

$$\phi_{t|z_{t-1}=1} - \phi_{t-1} = -\frac{(1 - \phi_{t-1})\phi_{t-1}}{\chi_0/(\chi_0 - \chi_1) - \phi_{t-1}},$$

to unity), her beliefs are “entrenched” and relatively unresponsive to the current abuse realisation.¹⁸ Our modelling approach will thus be to avoid specifying any particular model of male decision-making, preferences and costs, but to instead directly estimate the potential role of self-control behaviour by treating $\underline{\phi}$, $\bar{\phi}$ and ι as parameters to be estimated.

Using the extreme value distribution of the emotional cues, we thus specify the risk of abuse from a violent partner as,

$$\chi_0(k_t, \bar{w}_t^2, \bar{w}_t^h, n_t, t, \phi_t) = \iota(\phi_t) \times \frac{\exp(\varrho(k_t, \bar{w}_t^2, \bar{w}_t^h, n_t, t))}{1 - \exp(\varrho(k_t, \bar{w}_t^2, \bar{w}_t^h, n_t, t))}, \quad (13)$$

where $\iota(\phi_t) = \iota \in [0, 1]$ if $\phi_t \in [\underline{\phi}, \bar{\phi}]$ and $\iota(\phi_t) = 1$ otherwise. This flexible specification can thus be argued to accommodate several alternative theories of abuse and several possible scenarios with regards to perpetrator self-control behaviour. The women in the model are assumed to be aware of risk-of-abuse function specified in (13) and take it into account both directly in relation to their choices (e.g. the effect of the labour supply choices) and in relation to their belief updating.

3.2.4 Marriage, Beliefs and Choices

The marriage and learning side of the model follows the simplified version above. A woman who enters period t as married can choose to remain married or divorce. A single woman meets a new prospective partner with probability $\varsigma \in (0, 1)$. The fraction of encountered men who are of the non-violent type is ϕ_b^q , where the superscript indicates that we allow the male type distribution to depend on the woman’s level of qualification, $q \in \{0, 1, 2\}$. This implicitly assumes a woman’s qualification is the only relevant characteristics in terms of determining her likelihood of matching with a violent male, and also that there is no serial correlation in the type of men that women match with.¹⁹

Abuse $z_t \in \{0, 1\}$ is realised after the woman has decided on her current level of labour supply k_t and made her current conception decision $f_t \in \{0, 1\}$. Hence these decisions under uncertainty about potential exposure to abuse. If a woman decides to become pregnant at time t , she will give birth before the start of the following period. Thus letting n_t denote her number of children, we have that

$$n_{t+1} = n_t + f_t. \quad (14)$$

where $1 > \chi_0 > \chi_1 > 0$ guarantees that the expression is negative, that is that her beliefs deteriorate after experiencing abuse. The change in beliefs limits to zero when ϕ_{t-1} either goes to zero or to unity. Moreover it is easy to see that the change in beliefs is a convex function of the initial beliefs ϕ_{t-1} and thus has convex lower contour sets.

¹⁸In general the interval boundaries, $\underline{\phi}$ and $\bar{\phi}$, could depend on all state variables. For instance if having children reduces the wife’s likelihood of choosing divorce, then in principle, the interval $[\underline{\phi}, \bar{\phi}]$ could more narrow. However, as noted below, we generally find little evidence for any significant role played by self-control. For this reason we have not explored the potential role of other state variables.

¹⁹We have further explored whether there is any difference in the future incidence of abuse experienced by currently single women depending on whether or not they experienced abuse in the past. Overall we find no such difference. However, decomposing this result, we find that currently single women with a history of abuse are less likely to match with new partners, but are slightly more likely to experience abuse again if they do. The latter result is however based on only 65 women and not statistically significant.

The direct utility from children and conception in (4) is specified as

$$\Psi_t^n = \beta_1^n n_t - \beta_2^n n_t^2 + f_t \varepsilon_t^f, \quad (15)$$

where ε_t^f is a temporary utility shock from conceiving a child, assumed to be normally distributed with zero mean and variance σ_f^2 . As in the simple model we assume that the (direct) utility of marriage has a deterministic and a stochastic part so that

$$\Psi_t^m = \psi^m + \varepsilon_t^m, \quad (16)$$

where ε_t^m is normally distributed with zero mean and variance σ_m^2 . The random utility can be interpreted as a temporary match quality shock. The utility shocks ε_t^f and ε_t^m are assumed to be independent of the earnings shocks and independent of each other.

A woman's beliefs ϕ_t are updated exactly as in (1) and (2) while taking into account that χ_0 is given by (13). The expected disutility from abuse for a married woman with current belief ϕ_t in (4) is given by $\bar{\Psi}_t^z = \pi(k_t, \bar{w}_t^2, \bar{w}_t^h, n_t, t, \phi_t) \psi^z$ where

$$\pi(k_t, \bar{w}_t^2, \bar{w}_t^h, n_t, t, \phi_t) = \phi_t \chi_1 + (1 - \phi_t) \chi_0(k_t, \bar{w}_t^2, \bar{w}_t^h, n_t, t, \phi_t), \quad (17)$$

is her perceived probability of experiencing abuse in the current period, and where

$$\psi^z = \psi_0^z + \psi_a^z a, \quad (18)$$

is the direct disutility of abuse. In order to allow for potential heterogeneity in “tolerance” of abuse, we allow in (18) for the possibility that high productivity type women ($a = 1$) have a different disutility of abuse compared with low productivity type women ($a = 0$).²⁰

4 Estimation

The model is estimated using the method of simulated moments (McFadden, 1989; Pakes and Polard, 1989). This approach entails, for any trial parameters, first solving the model using backwards induction. In doing this we are using a full numerical solution method, solving the *E*max function at every $t = 1, \dots, T$ (Keane and Wolpin, 1994).²¹ The deterministic part of state space at time t

²⁰Note also a woman's disutility from abuse, specified in (18), does not depend on the presence or number of children. This restriction mainly reflects that it is difficult to identify any such effect using the current data. Attempts at including the number of children as a determinant of the disutility of abuse has led fertility to be highly sensitive to this parameter.

²¹The *E*max function at time t is the expected value of the maximum over alternative-specific value functions at time $t + 1$ for any given point in the state space corresponding to each possible alternative in the choice set at time t .

is $\{n_t, \phi_t, x_t, m_{t-1}, k_{t-1}, t, q, a\}$. After solving, the model is then forward-simulated to obtain simulated panel data with lifecycle paths for a large number of individuals with a distribution of observable characteristics that correspond to those observed in the data.

4.1 Simulated Population and Sampling

For any trial parameters outcomes are simulated for 15,000 women with a distribution of academic qualifications – the only source of observed initial heterogeneity – as observed in the data. When computing the simulated moments we focus on outcomes between the ages 17 to 40 to help correct for the initial conditions problem and end-of-horizon effects.

To account for the choice-based sampling frame used by the ALSPAC, we adopt a corresponding sampling frame on our simulated data. In particular, when computing the matched moments on the simulated data, we include every first birth from the moment of conception along with the following six periods for that woman. This places us as close as possible to the timing of the ALSPAC sampling frame, where women are first observed a few months into the study pregnancy.²²

4.2 Identification

Overall, 51 parameters are estimated using 101 empirical moments that are both static and dynamic in nature. Here we discuss how the model parameters are identified. In Appendix C we give examples of estimates of the sensitivity of the estimated parameters with respect to the empirical moments, corresponding to those shown in Figure 1, in order to highlight how our estimated abuse-related parameters in particular respond to various moments.

In the simple version of the model above, we argued that the type frequency ϕ_b and the type-specific behaviours, χ_1 and χ_0 , were identified from the persistence of abuse, the abuse risk in the early stages of a relationship, and the overall level of abuse. The same logic continues to apply in the full model. More generally, with longitudinal data, the binary male type distribution will be strongly reflected in the count distribution as well as in the persistence of abuse.²³ Also as in the simple version, the overall divorce probability and the divorce probability specifically after abuse help identify the size of the marital utility shocks and the disutility of abuse.²⁴ Note that the rate at which women selectively divorce abusive males

²²Standard errors are obtained by taking the square root of the diagonal elements of the variance-covariance matrix $\mathbf{Q} = (1 + \frac{1}{N}) (\mathbf{S}'\mathbf{W}\mathbf{S})^{-1}$ where \mathbf{S} is the first derivatives of the vector of moments \mathbf{s} with respect to the parameter vector θ and N is the number of simulations. \mathbf{S} is numerically approximated using parameter bump sizes that vary between .01% and 1% depending on the sensitivity of the moments. For weighting matrix \mathbf{W} we use the inverse of the variances of the data moments.

²³For instance, using the count distribution of abuse in our seven periods of data, a standard χ^2 goodness-of-fit test can be used to reject that the observed count distribution is generated by a binomial process where abuse is i.i.d. over all women and periods. It would predict too few zeros occurrences and too few three or more occurrences. Moreover, such a process would have a low persistence of abuse. Instead, a two-type binomial mixture model can replicate the count distribution and the persistence well.

²⁴Interestingly, the observed rates of abuse help identify the marital utility shock, which has been difficult to identify in discrete choice dynamic programming models that do not incorporate domestic abuse data (see e.g., Keane and Wolpin (2010) and Sauer (2015)).

will also be strongly identified by the negative empirical relationship between abuse risk and partnership duration (highlighted in Appendix A, Figure A2).

A key addition in the full model is the modelling of fertility. The parameters of the utility from children and the fertility utility shock will naturally be identified by the distribution of the number of children and the fertility rates e.g. by marital status and by age. The observed delay in fertility within marriage will also reflect the learning environment and hence help identify the latter.²⁵

In order to identify the parameters of the gender-tension function (12), we match how abuse risk varies with labour supply, age, and earnings and the fertility rate by abuse status. In the full model we allow the male type frequency to vary with the woman's qualification level and to identify this variation we match abuse rates by qualification. How abuse varies by duration of partnership will inform the estimates of the parameters related to self-control, in particular it's (low) efficiency.

The remaining set of moments included in estimation can be broadly split into two main groups by what they help to identify. The first group contains moments related to employment (employment transitions and employment status by age, marital status, and qualifications) and wages (by employment level and qualification level, and for husbands). These moments strongly identify the parameters associated with the wage offer functions, the unobserved ability structure, the disutility of work effort, income associated with non-employment, and the correlation between per-period earnings shocks. The identified earnings structure combined with the observed marriage rate further identifies the sharing parameter. The second main group of empirical moments relates to fertility, including average age and average partnership duration at first birth, the distribution of completed fertility, the proportion of out-of-wedlock births, and birth rates for single and married women. These moments help identify the utility of children, conception utility shocks, child-related costs, and the level of income-support for lone mothers.²⁶

The discount factor and the parameter of relative risk aversion are not estimated but rather fixed at levels consistent with previous literature. The discount factor δ is set at 0.95 and the parameter of relative risk aversion λ is set at 0.7. Identification of δ and λ is a common problem in discrete choice dynamic programming models.

5 Estimation Results

In this section, we report the results from the estimation of the full model presented in Section 3. We first present the moments included in the estimation and the estimated model's fit to these moments. We

²⁵It should be noted that the current model emphasises the uncertainty about the husband's type as the source of delay in fertility within marriage at the expense of other potential sources of uncertainty, e.g. about some underlying match quality as in Brien *et al.* (2006) or about the husband's earnings capacity. Also, the fact that wages are not directly observed but imputed using occupational information means that we are not able to capture persistence of earnings shocks.

²⁶As an auxiliary moment we include the fraction of women who remain childless. As this empirical moment, per construction, cannot be computed in the ALSPAC data, we obtain it from Figure 2 in ONS (2013a). We similarly obtain the average age at first marriage for the relevant cohorts from ONS (2013b).

then present the fit to a further set of moments, including moments relating to the association between children and abuse. Finally, we present and discuss the parameter estimates.

5.1 Moments and Model Fit

5.1.1 Fitted Moments

A comparison of empirical and simulated moment values are presented in Table 2 below along with Tables B1 - B3 in Appendix B. Table 2 presents abuse-related moments. Table B1, B2, and B3 presents moments related to marriage/fertility (marital transitions, birth rates by marital status, completed fertility, age at first birth and marriage, average duration of partnership at first birth and at divorce), employment (labour supply status by age, qualification level and marital status, and employment transitions), and hourly wages (by labour supply status, of husbands, and by qualification level and age respectively).

We focus here particularly on how the model fits a set of core abuse-related moments. Panel A of Table 2 shows that, in line with the simple illustrative model above, the full empirical model replicates quite closely the overall level of abuse and the abuse transitions. Closely related, the model predicts very well also the count distribution of abuse incidents (Panel B) over the seven periods. Importantly for the identification of the type-structure and selective divorce behaviour, Panel C shows that the model captures well how the incidence of abuse varies with partnership duration.²⁷

Panel D shows that the model somewhat over-predicts the qualification gradient in abuse. It should be noted that the model predicts that high qualified women experience a markedly lower rate of abuse even though the parameter estimates do not suggest that they meet violent men at a particularly lower rate (see below). Instead, the lower incidence of abuse experienced by high qualified women reflects that they are more frequently working, have fewer children and have them later.

Panel E shows that the model captures that abuse declines with age, though it somewhat under-predicts the particularly high abuse incidence among young mothers. A concern with the current data is that, while it is very rich in terms of characterising the level and dynamics of abuse experienced by mothers, it does not provide any corresponding information for non-mothers. In order to tackle this concern, we use information from the 1996 British Crime Survey (ONS, 1996). This was the first year in which the BCS piloted a computer-assisted self-completion module on domestic violence and the particular year of course also falls within the sample period for our main ALSPAC data. From the BCS data, we include information on abuse experienced by childless women (Mirrlees-Black, 1999) which shows that

²⁷As the simple illustrative model above highlighted, the abuse rate early in relationships is a key moment with several parameter estimates being quite sensitive to it. A potential concern is then whether the early abuse rate reported in Table 2 also holds in the general population. As it turns out, the finding is indeed consistent with evidence from the British Crime Survey (see below). Reporting on findings from the 2001 BCS, Walby and Allen (2004) note (p. 63 and Table 4.4 on p. 64) that “Domestic violence starts early in the life of those relationships where it is present. If domestic violence was going to become a repeated act, it had started during the first year of a relationship for 49 percent of women”. As we will see, our estimates suggest that the probability of matching with a violent type (i.e. a repeat abuse perpetrator) is about 20 percent whereby our unconditional estimate of the first-year abuse rate experienced by the ALSPAC first-time mothers in Panel C of Table 2 is well in line with the conditional rate reported by Walby and Allen for the general population.

Table 2: Matched moments: abuse.

Panel A: Abuse Rate and Abuse Transitions			
	Mean	No Abuse at $t + 1$	Abuse at $t + 1$
No abuse at t	0.922 <i>0.921</i>	0.947 <i>0.954</i>	0.053 <i>0.046</i>
Abuse at t	0.078 <i>0.079</i>	0.543 <i>0.592</i>	0.457 <i>0.408</i>
Panel B: Count Distribution of Abuse Incidents			
0	1-2	3-4	5+
0.764 <i>0.709</i>	0.163 <i>0.239</i>	0.054 <i>0.023</i>	0.019 <i>0.029</i>
Panel C: Abuse Rate by Partnership Duration in Years			
0-1	2-3	4-5	7+
0.133 <i>0.129</i>	0.088 <i>0.108</i>	0.068 <i>0.090</i>	0.070 <i>0.055</i>
Panel D: Abuse Rate By Qualification Level			
Low Qual.	Medium Qual.	High Qual.	
0.085 <i>0.101</i>	0.078 <i>0.068</i>	0.075 <i>0.074</i>	
Panel E: Abuse Rate By Age Group			
Age 17-24	Age 25-32	Age 33-40	
Mothers			
0.117 <i>0.091</i>	0.068 <i>0.068</i>	0.078 <i>0.091</i>	
Childless Women			
0.110 <i>0.061</i>	0.090 <i>0.091</i>	0.071 <i>0.099</i>	
Panel F: Abuse Rate By Labour Supply at $t - 1$			
Not Working	Part-Time	Full-Time	
0.084 <i>0.079</i>	0.072 <i>0.065</i>	0.093 <i>0.096</i>	
Panel G: Abuse Rate By Potential Income Quartile			
Q1	Q2	Q3	Q4
0.090 <i>0.090</i>	0.078 <i>0.095</i>	0.075 <i>0.061</i>	0.069 <i>0.072</i>
Panel H: Divorce and Birth Rate by Abuse Status at $t - 1$			
Divorce Rate if		Birth Rate if	
Non-Abused	Abused	Non-Abused	Abused
0.014 <i>0.020</i>	0.076 <i>0.087</i>	0.202 <i>0.222</i>	0.124 <i>0.118</i>

Notes: All empirical moments are for the ALSPAC sample of 3,960 first-time mothers presented in detail in Appendix A. This excepts the abuse rate for childless women which comes from the 1996 British Crime Survey (See Mirrlees-Black, 1999). The numbers in italics are the corresponding moments in the estimated baseline model.

the age-profile of abuse for childless women is similar to that for mothers.²⁸ We include these moments in

²⁸The BCS survey is a repeated cross-sectional survey and as such does not provide any longitudinal information that could be used to study e.g. persistence of abuse from one year to the next.

the estimation, and the table shows that we match age-group-specific abuse rates reasonably well, though the model again somewhat underestimates the rate of abuse experience by the youngest age group.

The model further replicates the U-shaped relationship between labour supply and exposure to abuse (Panel F), implying that part-time work is the labour supply status least associated with abuse. As will be seen below, the estimated abuse-probability function indicate very little direct difference between part- and full-time work. Instead, the observed difference is explained by endogenous labour supply choices: part-time work tends to be chosen by women with more positive beliefs about their partners' nature, with longer partnership duration, and with a larger number of children. In contrast, full-time work is more commonly chosen by women with less positive expectations about their partners' nature, with shorter marriage duration and with fewer children.

Panel G highlights the relationship between potential earnings (hourly full-time wage) quartile and incidence of abuse, showing a clear negative gradient. The estimated model exhibits a similar, though somewhat steeper, gradient. Finally, Panel H shows that the model predicts well that women who experience abuse at time t are substantially – about five times – more likely to divorce in the following period, and also substantially less likely to conceive a further child.²⁹

The fit to all other moments used in the estimation are presented in Tables B1 to B3 in Appendix B. Table B1 shows that the model fits the marital transitions well. The model slightly over-predicts births to single women but predicts quite well the proportion of women who remain childless and the distribution of number of children among those who do have children. Importantly, the model predicts the timing of first births very well, both in terms of the mother's age and in terms of partnership duration. It also replicates fairly accurately the average duration at divorce.

Table B2 shows the model's fit to labour supply moments. The table shows that the model replicates a key set of stylised facts well: (i) the majority of transitions into employment in the specific population of mothers with young children are into part-time employment; (ii) younger mothers are the least likely to work whereas older mothers are the most likely to work part-time; (iii) for both single and married mothers, not working is the most common labour supply status; (iii) more qualified mothers work more than less qualified mothers. Table B3 shows that the model correctly predicts that the accepted wages of full-time workers exceed those of part-time workers. The model also predicts a realistic qualification gradient for accepted hourly wages.

5.1.2 Additional Moments

As an important check we also consider the fit to a number of unmatched moments. Consider first the relationship between abuse and children. In our model specification, we allowed the “tension” factor $\varrho(\cdot)$ to depend on the number of children. However as we will show below, the estimated model suggests only

²⁹Reflecting the importance of these moments, we match in addition to the levels, also the differences by abuse status. The empirical annual birth rates are for the periods following the birth of the ALSPAC child and hence capture births of subsequent siblings, and the simulated birth rates are computed in the corresponding way.

a very modest direct effect of children on tension and hence on abuse risk. That of course does not mean that the incidence of abuse is in any way unrelated to children.

Indeed, an interesting feature is how the incidence of abuse evolves around child-birth. Panel (a) of Figure 3 plots the incidence of abuse over time for first time mothers starting in pregnancy ($Time = 0$). In terms of the ALSPAC data, the initial observation is based on the reports by first time mothers from the early stages of their pregnancies with the 12-month period in question thus relating mainly to abuse experience pre-pregnancy. In contrast, the first post-birth reported abuse incidence relates roughly to the second half of the pregnancy and the first eight months post birth. The increased rate of reported abuse during pregnancy and after birth compared to pre-pregnancy has received attention in the literature (Jasinski, 2004; Bowen *et al.*, 2005). According to the ALSPAC data, the rate of abuse nearly doubles between pre/early pregnancy and late pregnancy/child-birth.³⁰ After the birth the rate of abuse remains high, but reduces slightly over time.

The figure then compares the simulated data to this empirical pattern. Interestingly, without children having any marked direct estimated effect on abuse risk, the model replicates the incidence of abuse around child-birth remarkably well, predicting a marked increase between pre- and post-birth along with a gradual decrease thereafter. The main driver behind the predicted increase in abuse between pre- and post-birth is however not a direct effect of children, but rather a reverse causality relating to the underlying learning structure. Women in the model choose to conceive in response to positive beliefs about the partner's nature induced by the absence of abuse. Stated differently, pregnancies and births do not *per se* significantly increase women's exposure to abuse – rather pregnancies occur selectively as a response to an initial absence of abuse.

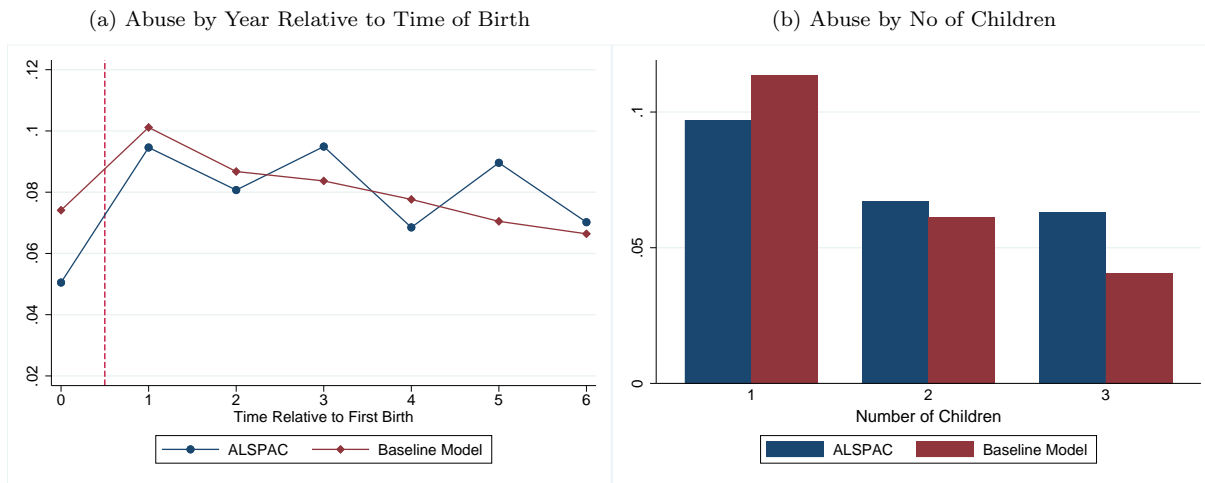
Panel (b) of figure 3 shows that the estimated model more generally predicts the relationship between the number of children and abuse incidence very well, with both the empirical and simulated data showing a decreasing relationship. Again, this relationship is best understood from the perspective of endogenous fertility: higher order births are more likely to occur in the absence of abuse.

The estimated model further left the relationship between marriage and qualifications unmatched. Panel A of Table B4 in Appendix B shows the proportion married by qualification level. There is a marked positive gradient in the data and a similar gradient occurs in the simulated model, reflecting a stronger financial incentive for more qualified women to marry under the assumed assortative mating. The table also reports (Panel B) the empirical and predicted divorce rates, both following abuse and non-abuse, by qualification level. In the empirical data, the divorce rate decreases with the wife's qualification level, both following abuse and non-abuse, and part of these qualification gradients are reflected also in the simulated data.

More directly related to the question of whether the presence of children make it harder to leave an abusive relationship, Panels C and D of Table B4 present the empirical and predicted divorce rates by

³⁰The abuse rate in pregnancy is included as a matched moment in the estimation whilst the post-birth abuse rates are not.

Figure 3: Abuse by year relative to first birth and by number of children.



Notes: The empirical moments presented in the figures are for the ALSPAC sample of 3,960 first-time mothers presented in detail in Appendix A. The figures also present the corresponding moments in the estimated baseline model.

number of children and by years since first birth. Panel C shows that the divorce rate is lower in the presence of a larger number of children. This holds both following abuse and non-abuse and both in the empirical and the simulated data. Panel D shows that, in the empirical data, in the absence of abuse the divorce rate is seemingly constant with respect to time since first birth. In contrast, the divorce rate following abuse is markedly higher in early after first birth compared to later on. These patterns are also replicated by the model.

Finally, related to role of learning, Panel E of Table B4 present the empirical and predicted divorce rates, after the birth of the first child, by the duration of marriage at the time of birth. This shows that, in the absence of any abuse, the divorce rate is shows no specific relation to marriage duration. In contrast, and consistent with learning, the rate of divorce in response to abuse is markedly higher for marriages that had short durations at the time of birth. Both these patterns are replicated by the estimated model.

5.2 Parameter Estimates

The estimated parameters are reported in Tables 3 and 4. Table 3 presents the estimated β -coefficients from the wage-offer equations (6), (9) and (10), the ability-qualification equation (7), and the child-related equations (11) and (15). More specifically to the current model, Table 4 reports all remaining parameters relating to marriage, fertility and abuse. We will briefly comment on the estimates in Table 3 before providing a more in-depth discussion of the parameters in 4.

Table 3: Parameter estimates: linear equations.

Panel A: Wage Offer Functions				
	Non-Emp. w^0	PT Emp. $\log(w_t^1)$	FT Emp. $\log(w_t^2)$	Husband $\log(w_t^h)$
Constant	3,115.6 (47.2)	6.854 (0.020)	7.391 (0.036)	9.404 (0.089)
a		0.904 (0.025)	1.017 (0.010)	0.256 (0.785)
x_t		0.090 (0.001)	0.090 (0.001)	
$x_t^2/100$		-0.147 (0.001)	-0.147 (0.001)	
age_t				0.016 (0.005)
$age_t^2/100$				-0.0004 (0.0048)
Panel B: Cholesky Terms				
a_{22}	a_{32}	a_{33}	a_{h2}	a_{hh}
0.040 (0.001)	0.127 (0.068)	0.206 (0.007)	0.290 (0.338)	-0.161 (0.325)
Panel C: Child-Utility, Childcare Costs and Income Support				
	Child Utility	Childcare Cost	Income Support Single Mothers	
n_t	1.367 (0.039)	5,659.1 (43.5)	2,477.8 (36.4)	
n_t^2	-0.0001 (0.0003)	-100.00 (327.36)	-825.92 (87.31)	
Panel D: Ability Probability Function				
Constant	0.338 (2.698)			
$q = 1$	0.475 (0.009)			
$q = 2$	0.518 (4.117)			

Notes: Parameters from baseline model estimated by simulated method of moments using empirical moments from the sample of first time mothers in the ALSPAC data. Standard errors in parenthesis. See Appendix A for details of sample.

Consider first the earnings regressions in Panel A of Table 3. The female earnings equations imply that high-ability women earn about 2.5 times as much as low ability women. The earnings growth due to an extra year's full-time experience ranges from about 20 percent at the early career stages down to zero percent growth for women who have worked full-time for fifteen years. The estimated maximum childcare costs presented in Panel C (incurred in full if working FT) are substantial, ranging from over to £5,000 per year with one child to over £16,000 with three children. The estimated child-related income available to single mothers is also substantial, starting at over £1,600 per year with one child.

Panel D presents the estimated relationship between the observable qualifications and the unobservable ability types. The probabilities of being high ability ($a = 1$) if low-, medium-, and high-qualified are 0.58, 0.69 and 0.70 respectively. Hence the low ability women are a minority group with a particular concentration among the low qualified.

Table 4: Parameter estimates continued: remaining parameters.

Panel A: Preference Parameters				
Marriage Utility		Fertility	Work Effort Cost	
ψ^m	σ_m^2	σ_f^2	μ_1	μ_2
639.29	1,732.1	3.939	0.997	0.981
(13.48)	(1,996.6)	(0.099)	(0.002)	(0.003)
Panel B: Abuse Parameters - Types/Disutility				
Baseline Type Frequencies			Abuse Disutility	
$\phi_b^{q=0}$	$\phi_b^{q=1}$	$\phi_b^{q=2}$	ψ^z	ψ_a^z
0.731	0.878	0.860	335.06	219.5
(0.004)	(0.001)	(0.019)	(5.28)	(569.3)
Panel C: Abuse Parameters - Abuse Freq.				
Non-Violent Type		Violent Type: Labour Supply		
χ_1		χ_0^0	χ_0^1	χ_0^2
0.043		1.038	0.376	0.340
(0.001)		(0.006)	(0.003)	(0.007)
Violent Type: Potential Earnings, Children and Age				
χ_0^w	χ_0^h	χ_0^n	χ_0^{n2}	χ_0^t
-0.133	0.011	-0.028	0.002	0.0003
(0.001)	(0.001)	(0.003)	(0.012)	(0.0047)
Panel D: Self-Control Parameters				
Trigger Feed-through		Belief Range		
ι		$\underline{\phi}$	$\bar{\phi}$	
0.972		0.581	0.907	
(0.506)		(0.025)	(0.102)	
Panel E: Sharing, Cost Fractions, Meeting Rate				
Sharing		Meeting Pr. Childcare		
τ	ς	ρ^0	ρ^1	
0.721	0.138	0.193	0.222	
(0.007)	(0.001)	(0.020)	(0.006)	

Notes: Parameters from baseline model estimated by simulated method of moments using empirical moments from the sample of first time mothers in the ALSPAC data. Standard errors in parenthesis. See Appendix A for details of sample.

Consider now the estimates in Table 4. Comparing the systematic utility from marriage ψ^m and the disutility from abuse ψ^z shows that, in line with the simple model above, abuse is associated with a large-scale negative utility shock – close to half of ψ^m (for low ability women and, adding ψ_a^z , larger for high ability women). The marriage utility shocks are large, with the variance σ_m^2 being around 2.5 times the size of ψ^m , again, similar to the simple model. The estimates of ϕ_b^q -values by qualification level indicate the low-qualified women are more likely to encounter violent-type men when encountering new

potential partners. The estimated abuse probability for non-violent males, χ_1 , is, as in the simple model above, naturally low.

Consider then the estimated parameters of the tension factor $\rho(\cdot)$ (12) that underlies the abuse probability $\chi_0(\cdot)$ for the violent type (13). In order to interpret the implied effects, note that the probability of abuse from a violent partner for a woman who is currently not working, with average (accepted) full-time earnings capacity, average husband earnings capacity, one child and aged 27, and a partner who is not exercising self-control (that is, with trigger “pass-through” rate $\iota = 1$) is $\chi_0(\cdot) = 0.77$. This is only slightly higher than the calibrated value of χ_0 in the simple version above.

Table 5 then highlights how the abuse risk $\chi_0(\cdot)$ is affected by labour supply, own and husband’s earnings capacity and with children. The estimates indicate a significantly higher risk of abuse for women who are not working: working part-time reduces the rate of abuse by a violent partner by close to 14 percentage points compared to not working. Further increasing her labour supply to full-time does not substantially additionally reduce the risk of abuse. Varying the wife’s earnings capacity from the 10th to the 90th percentile reduces the risk of abuse also by a sizeable 4.5 percentage points. In contrast, an increase in the husband’s earnings capacity is estimated to have only a negligible effect.

The estimates thus suggest both a direct abuse-reducing effect of labour supply and of the wife’s potential earnings. Whilst the latter effect is in line with the best causal evidence (Aizer, 2010), the former too has backing in the literature, most notably in relation to the US 1996 welfare reforms that imposed work requirements on benefit recipients. There was concern that this feature could potentially trigger a male backlash and increased domestic abuse. For this reason, the impact on abuse incidence was measured in a number of the localised randomised control trials that preceded the reform. Rather than finding an increase in abuse, these studies found either no effect or indeed a decrease, in some cases substantial (see Riger and Staggs, 2004 for a survey).³¹

Whilst we allow for the possibility of children directly affecting the abuse risk, the estimated direct effects are quite small. However, as seen as above these modest direct effects are still compatible with fairly sizeable associations between children – both in terms of births and numbers – and abuse incidence. Similarly, while we have seen, both in the actual and in the simulated data, that age and partnership duration are strongly negatively associated with abuse incidence (see Table 2), the estimated direct effect of time is also negligible. This indicates that these negative associations reflect the endogenous dynamics in relation to, for instance, partnership decisions, rather than any marked direct effect of time/age. Overall, these findings are more consistent with exposure and bargaining theory, and less with stress and identity-theory.

The estimated parameters indicate that, whilst violent males may exercise self-control over quite a sizeable range of beliefs for the wife, it does not defuse more than a small share (2.8 percent) of abuse

³¹Two of the most well-known randomised control trials in this context were the Minnesota Family Investment Program (MFIP) and the The National Evaluation of Welfare-to-Work Strategies (NEWWS). In both trials, women on the experimental welfare programmes experience reduced incidence of abuse compared to the control groups (Gibson *et al.*, 2003, Gennetian, 2003). For the MFIP in particular Gennetian and Miller, 2002 reported an 18 percent reduction.

triggers. Interestingly, the range of beliefs $[\underline{\phi}, \overline{\phi}]$ over which a violent-type man exercises self-control includes the baseline beliefs, ϕ_b^q , for all qualification levels. Hence, according to the model, violent men seek to exercise self-control at the initial stages of new relationships. However, as we have also seen, the rate of abuse is also empirically high at the start of relationships which explain why the estimated model suggest that self-control reduces abuse only marginally.

Table 5: Variation in abuse risk from violent partners.

Wife's Labour Supply	
Part-Time vs Not Working	Full-Time vs Part-Time
-0.137	-0.008
Earnings Capacity (Full-Time Wage)	
Wife 90% vs 10%	Husband 90% vs 10%
-0.052	0.002
Children	
1 Child vs No Child	2 Children vs 1 Child
0.020	0.004

Notes: The tables shows the estimated change in the probability of abuse from a violent ($r = 0$) type partner when (i) the wife's labour supply changes from not working to working part-time, and from working part-time to full-time, (ii) when her/his full-time earnings capacity changes from the 10th to the 90th percentile, and (iii) when the number of children changes from 0 to 1 and from 1 to 2. When not varied, the wife's labour supply is set to not working, earnings are set to mean values, and the number of children are set to one, and the husband's is assumed to not exercise self-control ($\iota(\phi_t) = 1$).

The estimated meeting rate ς is also effectively unchanged from the simple model. Conception utility shocks are important in the model, suggesting a fair amount of randomness in the timing of fertility: the estimated variance σ_f^2 is more than twice as large as the (annual) marginal utility of a child (see Panel B of Table 3). Childcare costs are reduced to about a fifth for women who are either not working or working part-time compare to working full-time. The estimated parameter τ indicates substantial income sharing.³²

6 Counterfactual Experiments

In this section we use the model to explore two distinct sets of questions. First, we explore the overall effect of uncertainty and learning on behaviour and outcomes. To do this we re-simulate the model under the counterfactual information structure where women can immediately observe any male's type as they meet. Second, we explore the effect of changes in the economic environment, focusing particularly on

³²It should be noted however that τ can also capture household public goods whereby the sum of her consumption as a proportion of total household income (τ) and his corresponding consumption as a proportion of total income can exceed unity.

aspects that economically “empower” women in general and mothers in particular. These experiments include (i) raising female wages to close the gender pay gap, (ii) increasing the child-related income available to single mothers, and (iii) providing subsidised child-care to households where the mother is working.

The simulations highlight how the interplay between labour supply and fertility in particular is key to the predicted impact of policy on the incidence of abuse with potentially substantial responses on both margins. In contrast, we find that marriage rates are much less responsive to policy. These findings are well in line with the existing empirical literature that has shown that fertility can be quite responsive to financial incentives generated by policy, but that the response crucially depends on the associated labour supply responses.³³ It should be borne in mind however that the literature on fertility responses to policy faces the general challenge of separating out responses that represent a shift in the *timing* of fertility from the longer run impact on *completed* fertility.

The empirical literature on the effects of financial incentives on marriage has largely used variation in marriage penalties or bonuses arising from changes to the tax-benefit code. While the estimated effects, if any, go in the expected direction, studies generally find that the effects on marriage are modest.³⁴

Whereas in the model estimation we focused on the population of first-time mothers in order to match the ALSPAC sample, the focus in this section is on the entire female population between the ages of 17 and 40. We do however also consider the incidence of abuse experienced by mothers and non-mothers respectively. This is of specific interest as a substantial literature argues that there are negative effects on children’s outcomes and behaviours of witnessing abuse between parents (McTavish *et al.*, 2016).

The results from the counterfactual simulations are presented in Table 6 and Figures 4 - 5. Table 6 presents results for a set of statistics computed across the women’s lifetimes. Figure 4 highlights some

³³Strong effects on fertility have generally been found for direct pro-natalist policies, a leading example being Milligan (2005) who studies the Allowance for Newborn Children introduced in Quebec in 1998. Substantial effects have also been found for tax-based fertility incentives, for instance by Laroque and Salanie (2014) for France and for publicly provided childcare by Bauernschuster *et al.* (2016). The estimated effects of in-work benefit reforms have on the other hand been more mixed, reflecting the fact that the positive fertility incentives are bundled with labour supply incentives, making the overall predicted effect ambiguous. For the US, Baughman and Dickert-Conlin (2009) found a very small negative effect on fertility of the Earned Income Tax Credit expansion. For the UK Francesconi and van der Klaauw (2007) found that the 1999 Working Families Tax Credit reform generated an increase in the labour supply of single mothers and a (statistically insignificant) reduction in their subsequent fertility. Similarly, for the same reform, Brewer *et al.* (2012) found no increase in fertility for single women but an increase in the fertility of coupled women for whom the labour supply incentives would have been weaker.

³⁴Key contributions based on US data include Dickert-Conlin and Houser (2002), Eissa and Hoynes (2000), Bitler *et al.* (2004), and Fisher (2013). For instance, Eissa and Hoynes (2000) find that a \$1,000/year increase in the marriage tax penalty reduces the married rate by 0.4 percent, whilst Fisher (2013) finds a larger effect at 1.7 percent. Studies using larger reforms have found mixed results. Again for the US Dickert-Conlin and Houser (2002) find little or no effect of expansion of the Earned Income Tax Credit on marriage, whereas Bitler *et al.* (2004) found the 1996 US welfare reform reduced both marriage- and divorce rates. Fisher and Zhu (2019) found that reduced welfare payments to lone parents in Australia increased the rate of repartnering after separations. UK-based studies are rare and have found mixed effects. Francesconi and van der Klaauw (2007) study the effect of the 1999 Working Families Tax Credit reform that increased in-work benefits to households with at least one child on a variety of outcomes. By effectively contrasting the rate of entry into partnership for single mothers to that of single women without children, they conclude that the WFTC reform had a small but significant negative effect on the partnership entry rate for single mothers. Using that the same reform, on average, marginally reduced the benefit-induced partnership penalties, Anderberg (2008) found that the WFTC reform marginally increased partnership rates, with response rates similar to those estimated by Fisher (2013) for the US.

more details of the dynamics of the responses by presenting various outcomes – relative to the baseline model – by age. Figure 5 focuses in particular on the timing of conceptions relative to first marriage.

Table 6: Counterfactual simulations: lifetime outcomes.

	Baseline Model	Perfect Information	Policy Experiments		
			Increased Female Wages	Income Support Single Mothers	Subsidized Childcare
Age at First Marriage	21.99	22.11	21.99	21.99	21.99
Divorce Rate	0.033	0.024	0.034	0.033	0.033
Age at First Birth	27.6	22.5	28.2	26.8	26.4
Proportion Childless	0.158	0.063	0.194	0.094	0.058
Average Nr of Children: All	1.86	2.16	1.81	2.01	2.14
Low Qualified	1.58	2.04	1.50	1.76	1.89
Medium Qualified	2.04	2.17	1.91	2.15	2.25
High Qualified	1.90	2.26	1.95	2.06	2.21
Non-Employed	0.358	0.384	0.161	0.409	0.337
Working Part-Time	0.182	0.336	0.209	0.190	0.222
Working Full-Time	0.460	0.280	0.630	0.401	0.441
Abuse Frequency: All	0.072	0.057	0.071	0.072	0.071
Low Qualified	0.100	0.078	0.100	0.100	0.100
Medium Qualified	0.058	0.047	0.057	0.058	0.058
High Qualified	0.063	0.050	0.062	0.063	0.063
Mothers	0.072	0.057	0.070	0.073	0.074
Non-Mothers	0.071	0.058	0.072	0.070	0.067

Notes: The table shows the outcomes from model simulations for a population aged 17-40 and with the same qualification distribution as observed in the ALSPAC sample (see Appendix A). The “baseline model” uses the parameter estimates from our preferred empirical specification as outlined above. In the “perfect information” case, the wife can directly observe the type $r \in \{0, 1\}$ of any man she encounters. For a description of the three policy experiments, see the text.

6.1 The Effect of Uncertainty

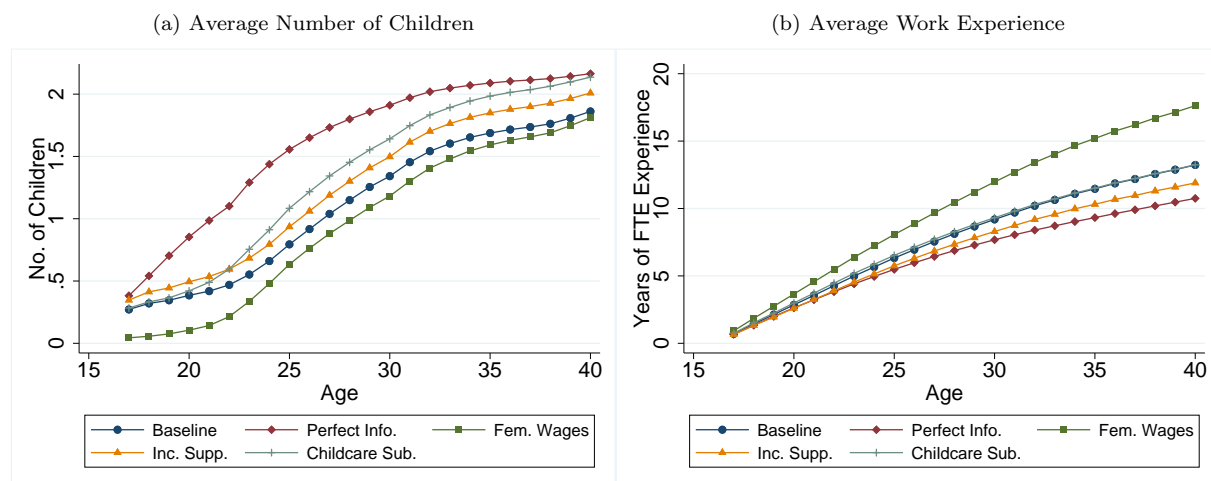
In the first counterfactual simulation we explore how uncertainty about males’ types affects women’s choices and outcomes. We focus here on the extreme opposite scenario relative to the baseline case, namely the case where any woman can immediately observe the type of any potential new partner.³⁵

There are two immediate behavioural consequences of the unobservability of a partner’s nature. First, when male types are not observable, women cannot directly reject marriage proposals from violent types. As that would be possible with perfect information, a lack of information increases the proportion of women who are married in early adulthood, generating a small decrease in the average age at first marriage and, more starkly, a large increase in the overall divorce rate (Table 6).

Second, uncertainty about a partner’s type affects fertility incentives. In particular, relative to perfect information, uncertainty about the partner’s type creates an incentive for delaying fertility within marriage in order to observe their behaviour. From the simple model above we know that learning occurs

³⁵We have further explored intermediate cases where a woman receives a binary signal $s \in \{0, 1\}$ which is correlated with the male’s true type, $\Pr(s = 1|r = 1) = (1 + \epsilon)/2$ and $\Pr(s = 1|r = 0) = (1 - \epsilon)/2$ for some value $\epsilon \in [0, 1]$. Based on the signal s she can then decide whether or not to marry this male. ϵ parameterises the precision of the signal with $\epsilon = 0$ corresponding to the baseline model (no information) and $\epsilon = 1$ the full information case. The results from these simulations indeed suggest that behaviour and outcomes with positive but imperfect information is, as expected, “between” the cases of no information and full information.

Figure 4: Average number of children and labour market experience in the counterfactual simulations.



Notes: The table shows the outcomes from model simulations for a population aged 17-40 and with the same qualification distribution as observed in the ALSPAC sample (see Appendix A). For details of the counterfactual scenarios, see notes to Table 6. Since, in the model, experience increases by $k_t \in \{0, 1, 2\}$, years of FTE experience is $x_t/2$.

quite fast, providing a strong incentive to delay fertility by a few years. This effect of uncertainty is highlighted in Figure 5 which plots the conception rate in years around the time of first marriage (where $year = 0$ indicates the year of first marriage). With perfect information, the conception rate is very high in the first three years of marriage and then drops substantially. In contrast, in the baseline model with uncertainty about partner's type, the conception rate is not only generally lower but, importantly, also no longer concentrated in the first few years of marriage. Uncertainty not only delays fertility, it also decreases overall fertility (Table 6), both in terms of increasing the proportion of women who remain childless and lowering the average number of children. This effect is particularly pronounced among low qualified women who, according to the estimated model are most likely to encounter violent-type men.

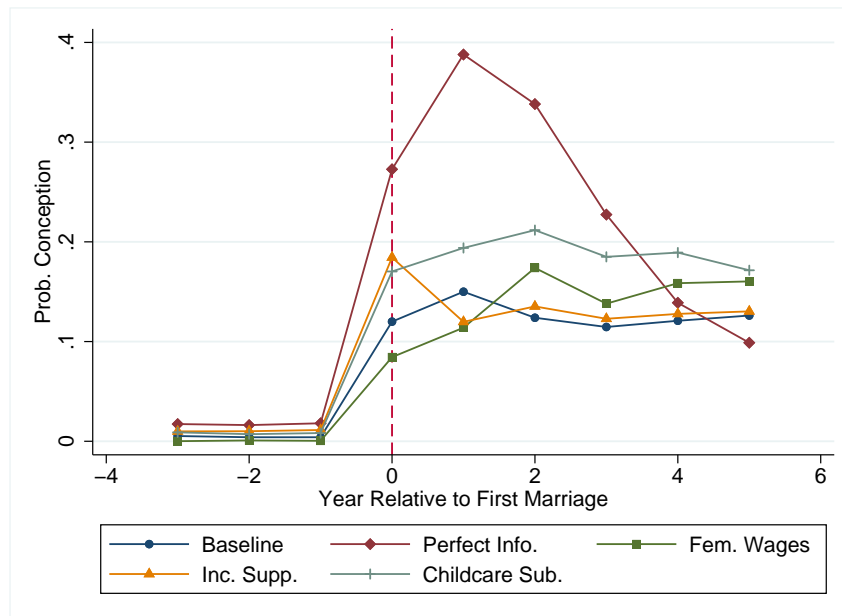
When male types are not observable women are naturally also more exposed to abuse. The overall abuse rate is 30 percent higher in the baseline model with uncertainty than in the perfect information scenario (Table 6). The increase in abuse incidence generated by uncertainty is particularly strong among mothers, reflecting that women more frequently have children with violent partners.³⁶

6.2 The Effects of Wages and Policy

We now revert back to the case where males' types are unobserved in order to focus on changes in the economic environment. Before highlighting differences between these cases two commonalities are worth

³⁶An interesting question is how a change in τ – the share of net household resources available to the wife when married (see eq. 5) – affects the incidence of abuse. A counterfactual simulation where τ is reduced to $\tau = 0.5$ results in women more frequently divorcing, less frequently being married overall, and more frequently working part-time, and experiencing about 2 percent less abuse overall.

Figure 5: Conception rate in years around year of first marriage in the counterfactual simulations.



Notes: The table shows the rate of conception by year relative to first marriage from simulations for a population aged 17-40 and with the same qualification distribution as observed in the ALSPAC sample (see Appendix A). For details of the counterfactual scenarios, see notes to Table 6.

noting. First, in all the cases considered, the impact on marriage rates is small – less than half a percentage point at every age. This should come as no surprise given that the literature has found married rates to be fairly unresponsive to financial incentives and given that none of the below experiments provide direct financial incentives for or against marriage.

Second, the simulated counterfactual scenarios share the feature from the baseline case that the rate of conception remains high for several years following marriage as opposed to peaking very early in marriage. This is consistent with learning about the nature of the partner before having children is of central importance.

6.2.1 Eliminating the Gender Wage Gap

In this counterfactual scenario we raise female earnings to the point where the average full-time earnings are the same for both genders. This involved a 15 percent increase. Part-time earnings were increased by the same proportion. In policy terms, this experiment could be thought of as representing a gender-specific wage subsidy.

Higher female wages encourage women to work more in the labour market. Labour market experience grows faster than in the baseline economy, and by age 40 the average experience is about a third higher than in the baseline economy (Figure 4). From Table 6 we also see that the average age at first birth increases by over half a year indicating an overall delay in fertility. Figure 4 shows that average achieved

fertility is lower at every age compared to the baseline case, and particularly so below the age of 25. Fertility is also delayed within marriage (Figure 5). Over time, fertility largely catches up so that, by age 40, the average number of children is less than three percent lower than in the baseline economy, with the decrease coming from low- and medium-qualified women.

Turning to abuse, we see that the overall incidence of abuse marginally lower, about two percent, than in the baseline scenario. This decrease in abuse incidence is driven by the increased labour force participation. The prediction that improved wages for women reduces exposure to abuse is in line with the findings in Aizer (2010), though the mechanism here is directly via increased labour supply rather than intra-household bargaining and the implied effect is also smaller.³⁷

Table 6 shows that the reduction in abuse incidence is concentrated among mothers. This reflects that fertility, by being delayed, is based on better information and is more selective. Hence, an important consequence of improved earnings opportunities for women is that children become less exposed to abuse between the parents.

6.2.2 Income Support for Single Mothers

The estimated model includes child-related income, $\beta_1^{ci} n_t + \beta_2^{ci} n_t^2$, available to single mothers, as a catch-all for either in- or out-of-work welfare benefits and potential child-support payments. At first glance, more generous income support to single mothers could potentially enable them to leave abusive relationships and could hence be a policy option for reducing domestic abuse. However more generous income support will also boost fertility incentives and lower labour supply through an expected income effect. Taking such broader responses into account, it is less clear that a generous child-related income support policy would indeed reduce the incidence of abuse. To explore this, we simulate the effect of an increase in the child support parameter, β_1^{ci} , by 20 percent relative to the baseline.

A first main effect of higher income to single mothers is to increase fertility by every measure: reducing the age at first birth, reducing the proportion who remain childless, and increasing the average number of children (Table 6). Age at first birth reduces by close to a year and, as can be seen from Figure 5, birth are somewhat more concentrated early in marriages.

As a result of having more children – and also due to the expected non-labour income effect – women work less. Furthermore, with less work experience pre-marriage, they are more likely to be out of the labour force when they eventually do get married. Given that being out of the labour force is associated with a higher rate of abuse from violent men, the reduced incentives for working indirectly increase exposure to abuse. Indeed, Table 6 indicates that increasing income to single mothers fails to reduce abuse incidence overall. Hence rather than reducing exposure to abuse, taking all behavioural responses into account – most notably fertility and labour supply responses – more generous income support to single mothers leaves women no less exposed to abuse. More worryingly, it leads to a higher rate of abuse

³⁷Aizer's estimates imply that a 15 percent increase in the relative wages of women would reduce women's exposure to assault by about 10 percent.

particularly towards mothers. This result is largely driven by the increase in early-in-marriage fertility, which implies that children are more frequently present during the critical early partnership stages.

6.2.3 Subsidised Childcare when the Mother is Working

The estimated childcare costs, $\beta_1^{cc}n_t + \beta_2^{cc}n_t^2$, apply equally to married and single mothers; however they are incurred in full only if the mother is working full-time, $\rho^2 = 1$, and partially at rate ρ^1 when she works part-time. Here we consider the effect of subsidised childcare for households with working mothers. To do so we reduce each fraction, ρ^2 and ρ^1 , of the full childcare cost incurred when the mother is working full- and part-time respectively by 20 percent.

Subsidised childcare has two main direct effects. First, it reduces the cost associated with working and hence encourages labour force participation among mothers. Second, it directly encourages fertility by reducing the overall expected cost associated with having children. As fertility and labour supply are negatively associated, this indirectly reduces labour supply, leaving the overall net effect on labour supply ambiguous.

The positive effect on fertility can be seen from Table 6: age at first birth decreases by over a year, the proportion who remain childless reduces substantially, and the average number of children increases including for the high-qualified women. Subsidised childcare leads to higher fertility in the early stages of marriage: note from Figure 5 that fertility in the first five years of marriage increases quite substantially relative to the baseline case.

The overall effect on labour supply is modest (Table 6). Consequently, the predicted impact on overall incidence of abuse is negligible, a conclusion that holds in each qualification category. Given that subsidised childcare is a popular policy option for simultaneously encouraging both fertility and labour supply, this would appear to be a positive conclusion, suggesting that such a policy can be used without increasing women’s exposure to abuse. However, the result comes with an important caveat: as can be seen from Table 6 the incidence of abuse among mothers – and hence the exposure to abuse of children – increases, in this case by about 3 percent.

7 Conclusions

Starting a relationship with a new intimate partner usually comes with hopes of a happy, long-lasting and well-functioning relationship. However, in far too many cases, such dreams fail to materialise as it is gradually disclosed that the new partner has a violent nature and will repeatedly engage in verbal and physical abuse. In formal modelling terms, this suggests that there is heterogeneity in male “violence types” which is not directly observable at the outset of a new partnership but is only revealed over time. Focusing on the impact of such uncertainty for women this paper has addressed two broad sets of questions.

First, what is the effect of uncertainty about a partner's violent nature on a woman's dynamic behaviour? For instance, does it lead to a delay in investments within marriage, most notably in fertility? Relatedly, what are the labour supply responses of women facing possible domestic violence? Do certain labour supply choices lead to an increased risk of abuse?

Second, what is the effect of female "economic empowerment" in the form of earnings opportunities and financial resources on the incidence of abuse? In particular, how do higher female wages affect women's choices and their exposure to abuse? What are the overall effects of better income support to single mothers and of subsidised childcare available to households in which the mother is working.

To address these questions, we constructed and estimated a dynamic lifecycle model where women meet and marry men, learn about their husbands' nature, and make decisions about fertility, labour supply, and about continued marriage or divorce. The core mechanism of the model is a learning process where a woman updates her beliefs about her husband's true nature by observing, over time, whether or not he engages in abusive behaviour. As the partner's type is gradually revealed, her perceived utility of continued marriage changes over time. But learning also indirectly affects fertility incentives. Children impose costs – either in the form of direct childcare costs or in terms of foregone earnings – which are shared whilst married. Hence, separating from a partner is more costly when children are present potentially trapping mothers in abusive relationships. Learning therefore implies an incentive for delayed child-bearing until more information is available about the partner's nature. It further affects labour supply decisions over time. A higher risk of divorce provides an incentive to build up labour market experience and earnings capacity in anticipation of potential singlehood. Moreover, in so far as some labour supply choices are more associated with abuse, a woman may avoid these particular choices early in relationships when the partner's nature is still largely unknown.

In order to study the various effects of uncertainty and learning on women's choices and outcomes, we used a counterfactual simulation of the model where a woman is provided with full information about the nature of any prospective new partner at the very moment they meet. In doing so, we uncovered some important interactions between learning and labour supply, marriage duration and fertility. Specifically, we found that, compared to the full-information scenario, the learning environment is associated with (i) earlier marriages, more frequent divorces, delayed fertility, and lower completed fertility, (ii) increased labour supply to avoid possible abuse and to build up labour market experience, and, of course, (iii) substantially higher rates of abuse. Arguably, the current model has focused on lack of information about the husband's type as the key source of uncertainty at the expense of other potential sources, e.g. learning about match quality or about earnings capacity. Allowing for alternative sources of uncertainty would be an interesting extension.

Counterfactual simulations were similarly used to analyse the effects of female economic empowerment in the form of higher female wages, increased income support provided to single mothers, and subsidised childcare when working. Higher female wages – to the point of eliminating the gender wage gap – were

unsurprisingly, found to increase female labour supply. It was also shown to reduce abuse incidence, although the implied effect was disappointingly small, a reduction of about 3 percent. Nevertheless, given the cost associated with domestic abuse – recently estimated at £66 billion per year for the UK (Oliver *et al.*, 2019), even such a relatively small decrease would have a substantial value (close to £2 billion per year).

Perhaps more surprising were the findings regarding more generous income support for single mothers and subsidised childcare. Such policies could, in principle, make mothers more financially independent and hence more able to walk away from abusive partners. However, we found that both policies also encourage fertility – most notably early in relationships – and the former policy in particular also decreases labour supply. With these policies, women more frequently find themselves in the early stages of relationships with children and with less accumulated labour market experience. As a result, they find it more difficult to leave abusive partners. Hence, we found that neither policy seems capable of reducing abuse towards women in general, and more worryingly, both policies can actually increase the incidence of abuse towards mothers in particular. The findings in relation to these policies in particular highlight the importance of jointly modelling fertility and labour supply choices.

The current model is the first to formally estimate a model where women learn the potentially abusive nature of their partners. To accomplish this, a set of assumptions have been imposed, including for instance rational Bayesian learning. Our model also does not incorporate any measure of health or well-being and does not consider any impacts on children beyond their existence. Hence there are many obvious directions in which this work could be extended.

Affiliations

¹Department of Economics, Royal Holloway University of London, Egham TW20 0EX, UK

²Management School, University of Liverpool, Liverpool L69 7ZH, UK

³CESifo, Munich, Germany

⁴Institute for Fiscal Studies, London

⁵IZA, Bonn, Germany

References

- Aizer, A. (2010). 'The gender wage gap and domestic violence', *American Economic Review*, vol. 100(4), pp. 1847–1859.
- Aizer, A. and Dal Bó, P. (2009). 'Love, hate and murder: Commitment devices in violent relationships', *Journal of Public Economics*, vol. 93(3), pp. 412–428.
- Alesina, A., Brioschi, B. and La Ferrara, E. (2021). 'Violence against women: A cross-cultural analysis for Africa', *Economica*, vol. 88(349), pp. 70–104.
- Amaral, S. and Bhalotra, S. (2017). 'Population sex ratios and violence against women: The long-run effects of sex selection in India', University of Essex.
- Anderberg, D. (2008). 'Tax credits, income support, and partnership decisions', *International Tax and Public Finance*, vol. 15, pp. 499–526.
- Anderberg, D. and Rainer, H. (2013). 'Economic abuse: A theory of intrahousehold sabotage', *Journal of Public Economics*, vol. 97(C), pp. 282–295.
- Anderberg, D., Rainer, H., Wadsworth, J. and Wilson, T. (2016). 'Unemployment and domestic violence: Theory and evidence', *Economic Journal*, vol. 126(597), pp. 1947–1979.
- Andrews, I., Gentzkow, M. and Shapiro, J.M. (2017). 'Measuring the sensitivity of parameter estimates to estimation moments', *Quarterly Journal of Economics*, vol. 132(4), pp. 1553–1592.
- Angelucci, M. (2008). 'Love on the rocks: Domestic violence and alcohol abuse in rural Mexico', *B E Journal of Economic Analysis & Policy*, vol. 8(1), p. Article 43.
- Bauernschuster, S., Hener, T. and Rainer, H. (2016). 'Children of a (policy) revolution: The introduction of universal child care and its effect on fertility', *Journal of the European Economic Association*, vol. 14(4), pp. 975–1005.
- Baughman, R. and Dickert-Conlin, S. (2009). 'The Earned Income Tax Credit and fertility', *Journal of Population Economics*, vol. 22(3), pp. 537–563.
- Berniell, I. and Facchini, G. (2021). 'COVID-19 lockdown and domestic violence: Evidence from internet-search behavior in 11 countries', *European Economic Review*, vol. 136(103775).
- Bhalotra, S., Britto, D.G., Pinotti, P. and Sampaio, B. (2021). 'Job displacement, unemployment benefits and domestic violence', CESifo Working Paper No. 9186.
- Bitler, M.P., Gelbach, J.B., Hoynes, H.W. and Zavodny, M. (2004). 'The impact of welfare reform on marriage and divorce', *Demography*, vol. 41(2), pp. 213–236.

- Bloch, F. and Rao, V. (2002). 'Terror as a bargaining instrument: A case study of dowry violence in rural India', *American Economic Review*, vol. 92(4), pp. 1029–1043.
- Bobonis, G.J., González-Brenes, M. and Castro, R. (2013). 'Public transfers and domestic violences: The roles of private information and spousal control', *American Economic Journal: Economic Policy*, vol. 5(1), pp. 179–205.
- Bowen, E., Heron, J., Waylen, A., Wolke, D. and ALSPAC Team (2005). 'Domestic violence risk during and after pregnancy: Findings from a British longitudinal study', *BJOG: An International Journal of Obstetrics & Gynaecology*, vol. 112(8), pp. 1083–1089.
- Bowlus, A. and Seitz, S. (2006). 'Domestic violence, employment, and divorce', *International Economic Review*, vol. 47(4), pp. 1113–1149.
- Boyd, A., Golding, J., Macleod, J., Lawlor, D.A., Fraser, A., Henderson, J., Molloy, L., Ness, A., Ring, S. and Smith, G.D. (2013). 'Cohort profile: The children of the 90s - the index offspring of the Avon longitudinal study of parents and children', *International Journal of Epidemiology*, vol. 42(1), pp. 111–127.
- Brewer, M., Ratcliffe, A. and Smith, S. (2012). 'Does welfare reform affect fertility? Evidence from the UK', *Journal of Population Economics*, vol. 25(1), pp. 245–266.
- Brien, M.J., Lillard, L.A. and Stern, S. (2006). 'Cohabitation, marriage, and divorce in a model of match quality', *International Economic Review*, vol. 47(2), pp. 451–494.
- Card, D. and Dahl, G. (2011). 'Family violence and football: The effect of unexpected emotional cues on violent behavior', *Quarterly Journal of Economics*, vol. 126(1), pp. 103–143.
- Coles, M.G. and Francesconi, M. (2019). 'Equilibrium search with multiple attributes and the impact of equal opportunities for women', *Journal of Political Economy*, vol. 127(1), pp. 138–162.
- Cosslett, S. (1981). 'Efficient estimation of discrete choice models', in (C. F. Manski and D. McFadden, eds.), *Structural Analysis of Discrete Data with Econometric Applications*, MIT Press.
- Dickert-Conlin, S. and Houser, S. (2002). 'EITC and marriage', *National Tax Journal*, vol. 55(1), pp. 25–40.
- Dugan, L., Rosenfeld, R. and Nagin, D.S. (2003). 'Exposure reduction or retaliation? The effects of domestic violence resources on intimate-partner homicide', *Law and Society Review*, vol. 37(1), pp. 169–198.
- Eissa, N. and Hoynes, H. (2000). 'Tax and transfer policy, and family formation: Marriage and cohabitation', University of California, Berkeley.

- Enander, V. and Holmberg, C. (2008). 'Why does she leave? The leaving process(es) of battered women', *Health Care for Women International*, vol. 29(3), pp. 200–226.
- Erten, B. and Keskin, P. (2018). 'For better or for worse? Education and the prevalence of domestic violence in Turkey', *American Economic Journal: Applied Economics*, vol. 10(1), pp. 64–105.
- Farmer, A. and Tiefenthaler, J. (1997). 'An economic analysis of domestic violence', *Review of Social Economy*, vol. 55, pp. 337–358.
- Fisher, H. (2013). 'The effect of marriage tax penalties and subsidies on marital status', *Fiscal Studies*, vol. 34(4), pp. 437–465.
- Fisher, H. and Zhu, A. (2019). 'The effect of changing financial incentives on repartnering', *Economic Journal*, vol. 129(623), pp. 2833–2866.
- Flatley, J., Kershaw, C., Smith, K., Chaplin, R. and Moon, D. (2010). 'Crime in England and Wales 2009/10', Home Office Statistical Bulletin.
- Francesconi, M. (2002). 'A joint dynamic model of fertility and work of married women', *Journal of Labor Economics*, vol. 20(2), pp. 336–380.
- Francesconi, M. and van der Klaauw, W. (2007). 'The socioeconomic consequences of "in-work" benefit reform for british lone mothers', *The Journal of Human Resources*, vol. 42(1), pp. 1–31.
- Garcia-Ramos, A. (2021). 'Divorce laws and intimate partner violence: Evidence from Mexico', *Journal of Development Economics*, vol. 150, p. 102623.
- Gemici, A. and Laufer, S. (2014). 'Marriage and cohabitation', Royal Holloway, University of London.
- Gennetian, A. (2003). 'Welfare policy and domestic abuse among single mothers', *Violence Against Women*, vol. 9, pp. 1171–1190.
- Gennetian, L.A. and Miller, C. (2002). 'Children and welfare reform: A view from an experimental welfare program in Minnesota', *Child Development*, vol. 73, pp. 601–620.
- Gibson, C.M., Magnuson, K., Gennetian, L., Duncan, G.J. and England, U. (2003). 'Does employment protect women from domestic violence?', Manpower Demonstration Research Corporation.
- González, L. and Rodríguez-Planas, N. (2021). 'Gender norms and intimate partner violence', IZA DP No. 14550.
- Gottfredson, M.R. and Hirschi, T. (1990). *A General Theory of Crime*, Stanford University Press.
- Guarnieri, E. and Rainer, H. (2021). 'Colonialism and female empowerment: A two-sided legacy', *Journal of Development Economics*, vol. 151, p. 102666.

- Haushofer, J., Ringdal, C., Shapiro, J.P. and Wang, X.Y. (2019). 'Income changes and intimate partner violence: Evidence from unconditional cash transfers in Kenya', NBER Working Paper 25627.
- Hidrobo, M. and Fernald, L. (2013). 'Cash transfers and domestic violence', *Journal of Health Economics*, vol. 32(1), pp. 304–319.
- Hsu, L.C. (2017). 'The timing of welfare payments and intimate partner violence', *Economic Inquiry*, vol. 55(2), pp. 1017–1031.
- Iyengar, R. (2009). 'Does the certainty of arrest reduce domestic violence? Evidence from mandatory and recommended arrest laws', *Journal of Public Economics*, vol. 93(1), pp. 85–98.
- Jasinski, J.L. (2004). 'Pregnancy and domestic violence: A review of the literature', *Trauma, Violence, and Abuse*, vol. 10(9), pp. 961–990.
- Keane, M.P. and Wolpin, K.I. (1994). 'The solution and estimation of discrete choice dynamic programming models by simulation and interpolation: Monte carlo evidence', *The Review of Economics and Statistics*, vol. 76(4), pp. 648–672.
- Keane, M.P. and Wolpin, K.I. (2010). 'The role of labor and marriage markets, preference heterogeneity and the welfare system in the life cycle decisions of black, hispanic and white women', *International Economic Review*, vol. 51(3), pp. 851–892.
- Laroque, G. and Salanie, B. (2014). 'Identifying the response of fertility to financial incentives', *Journal of Applied Econometrics*, vol. 29(2), pp. 314–332.
- Macmillan, R. and Gartner, R. (1999). 'When she brings home the bacon: Labor force participation and the risk of spousal violence against women', *Journal of Marriage and the Family*, vol. 61(4), pp. 947–958.
- Manski, C.F. and Lerman, S.R. (1977). 'The estimation of choice probabilities from choice based samples', *Econometrica*, vol. 45(8), pp. 1977–1988.
- McFadden, D. (1989). 'A method of simulated moments for estimation of discrete response models without numerical integration', *Econometrica*, vol. 57(5), pp. 995–1026.
- McTavish, J.R., MacGregor, J.C., Wathen, C.N. and MacMillan, H.L. (2016). 'Children's exposure to intimate partner violence: an overview', *International Review of Psychiatry*, vol. 28(5), pp. 504–518.
- Milligan, K. (2005). 'Subsidizing the stork: New evidence on tax incentives and fertility', *The Review of Economics and Statistics*, vol. 87(3), pp. 539–555.
- Mirrlees-Black, C. (1999). 'Domestic violence: Findings from a new British Crime Survey self-completion questionnaire', Home Office Research Study 191.

- Oliver, R., Alexander, B., Roe, S. and Wlasny, M. (2019). 'The economic and social costs of domestic abuse', Home Office Research Report 107.
- ONS (1996). 'British Crime Survey 1996', UK Data Service.
- ONS (1999). 'Labour Force Survey 1993-1999', UK Data Service.
- ONS (2013a). 'Cohort fertility: 2012', Statistical Bulletin.
- ONS (2013b). 'Marriages in England and Wales: 2010', Statistical Bulletin.
- Pakes, A. and Pollard, D. (1989). 'Simulation and the asymptotics of optimization estimators', *Econometrica*, vol. 57(5), pp. 1027–1057.
- Papageorge, N., Pauley, G., Cohen, M., Wilson, T., Hamilton, B. and Pollak, R. (2021). 'Health, human capital and domestic violence', *Journal of Human Resources*, vol. 56(4), pp. 997–1030.
- Piquero, A.R., Jennings, W.G., Jemison, E., Kaukinen, C. and Knaul, F.M. (2021). 'Domestic violence during the COVID-19 pandemic - Evidence from a systematic review and meta-analysis', *Journal of Criminal Justice*, vol. 74, p. 101806.
- Ramos, A. (2016). 'Household decision making with violence: Implications for conditional cash transfer programs', Universitat Autònoma de Barcelona.
- Riger, S. and Staggs, S. (2004). 'Welfare reform, domestic violence, and employment', *Violence Against Women*, vol. 10(9), pp. 961–990.
- Roberts, A.L., McLaughlin, K.A., Conron, K.J. and Koenen, K.C. (2011). 'Adulthood stressors, history of childhood adversity, and risk of perpetration of intimate partner violence', *American Journal of Preventive Medicine*, vol. 40(2), pp. 128–138.
- Sauer, R.M. (2015). 'Does it pay for women to volunteer?', *International Economic Review*, vol. 56(2), pp. 537–564.
- Schnor, C. (2015). 'Does waiting pay off for couples? partnership duration prior to household formation and union stability', *Demographic Research*, vol. 33, pp. 611–652.
- Stevenson, B. and Wolfers, J. (2006). 'Bargaining in the shadow of the law', *Quarterly Journal of Economics*, vol. 121(1), pp. 267–288.
- Tauchen, H.V., Witte, A.D. and Long, S.K. (1991). 'Violence in the family: A non-random affair', *International Economic Review*, vol. 32(2), pp. 491–511.
- Tertilt, M. and van den Berg, G. (2015). 'The association between own unemployment and violence victimizations among female youth', *Journal of Economics and Statistics*, vol. 235(4-5), pp. 499–513.

- Tur-Prats, A. (2019). 'Family types and intimate-partner violence: A historical perspective', *Review of Economics and Statistics*, vol. 101(5), pp. 878–891.
- Tur-Prats, A. (2021). 'Unemployment and intimate partner violence: A cultural approach', *Journal of Economic Behavior and Organization*, vol. 185, pp. 27–49.
- van der Klaauw, W. (1996). 'Female labour supply and marital status decisions: A life-cycle model', *The Review of Economic Studies*, vol. 63(2), pp. 199–235.
- Walby, S. and Allen, J. (2004). 'Domestic violence, sexual assault and stalking: Findings from the British Crime Survey', Home Office Research Study 276.
- WHO (2021). 'Violence against women prevalence estimates, 2018: global, regional and national prevalence estimates for intimate partner violence against women and global and regional prevalence estimates for non-partner sexual violence against women.', Geneva.
- Yau, N. (2019). 'The relationship timeline continues to stretch', <https://flowingdata.com/2019/03/19/>.