

The Dynamics of Domestic Violence: Learning about the Match*

DAN ANDERBERG^{a,b,c}, NOEMI MANTOVAN^d and ROBERT M. SAUER^{a,e}

a) Royal Holloway University of London; b) Institute for Fiscal Studies

c) CESifo, Munich; d) Bangor University; e) IZA, Bonn

October 22, 2017

Abstract

We present a dynamic lifecycle model where women choose partnership status, employment and fertility. Some males have a high propensity to engage in abusive behaviour, but women do not observe the nature of a prospective partner when they first meet. Instead, a woman learns her partner's type by observing his behaviour. This endogenous learning implies responses to abuse as women reassess the value of continuing and investing in their relationship. It further provides strategic incentives to delay relationship-specific investments within new relationships. The model is estimated by method of simulated moments using longitudinal data from the Avon Longitudinal Study of Parents and Children. We simulate alternative scenarios alternating, *inter alia*, the information that is available to women, the prevalence of abusive males, the gender wage gap, and the generosity of child support policy.

Keywords: Domestic violence, Learning, Fertility, ALSPAC

JEL Classification: J12, J13

*We are extremely grateful to all the families who took part in this study, the midwives for their help in recruiting them, and the whole ALSPAC team, which includes interviewers, computer and laboratory technicians, clerical workers, research scientists, volunteers, managers, receptionists and nurses. The UK Medical Research Council and Wellcome (Grant ref: 102215/2/13/2) and the University of Bristol provide core support for ALSPAC. This research was specifically funded by the British Academy (grant SG111103). We also thank seminar participants at CESifo (Munich), ISER (Essex), Oslo, Leicester, ESPE, and SEHO for helpful discussions and comments. We are especially grateful to Michael Keane and Ahu Gemici for comments that greatly improved the paper. Funding from the British Academy (grant SG111103) is gratefully acknowledged.

I Introduction

According to the latest estimates from the Crime Survey for England and Wales, over 8 percent of women experienced domestic abuse in 2014/15 (Woodhouse and Dempsey, 2016). According to the same survey domestic abuse accounts for about 20 percent of all violent incidents reported by respondents, and it has the highest rate of repeat victimization of any type of crime.

Economics has recently seen a surge in research on domestic violence which has provided a wealth of useful insights. This research has focused the role of labour market conditions, educational attainment, culture and social norms for the incidence of domestic violence (Aizer, 2010; Anderberg et al. 2015; Erten and Keskin, 2016; Alesina et al., 2016; Tur-Prats, 2016), on understanding triggers of domestic violence (Card and Dahl, 2011; Anderberg and Rainer, 2013), on the impact of policy – both related to law enforcement (Iyengar, 2009; Aizer and Dal-Bo, 2009) – and welfare and cash-transfers policy (Angelucci, 2008; Cormier, 2009; Bobonis et al, 2013; Ramos, 2016).

However, even with this recent stream of contributions, a number of core questions – particularly of 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; Tolman and Wang, 2005; and Riger and Staggs, 2004). This research has struggled with the notion 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) and on the notion of “reproductive coercion” exercised by abuse males (Clark et al., 2014). Finally, the perhaps most obvious – but also controversial – 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.

A key question when modeling partnership formation in particular concerns the information women have about the exact nature of their prospective partners. Do women upon entering partnerships know their partner’s abusive nature or is this something they learn over time through experience? In our model we incorporate learning in the simplest possible form. Any given man is assumed to either have a “violent nature” or a “non-violent nature”, and while his this determines his likelihood of being abusive, it is not directly observable to his wife. From the moment the partnership is formed, the wife will observe the behaviour of her partner and will update her beliefs based on her observations. Exposure to abuse gives rise to non-marginal changes in her beliefs, and hence in her expectations of what the future would hold within the relationship. This in turn triggers behavioural responses with respect to partnership status, labour supply, and child-bearing. But learning can furthermore have important effects on behaviour even prior to the incidence of any abuse. Intuitively, a woman may have a strategic incentive to delay relation-specific investments – most notably fertility – until she is reasonably certain that the partner does not have a violent nature.

One factor that has limited research on the dynamics of domestic violence has been the shortage of longitudinal data with representative populations and large sample sizes. We use the Avon Longitudinal Study of Parents and Children (ALSPAC), a local child-development survey that has followed a set of children from birth along with their parents. Our sample population will be the set of ALSPAC mothers. This implies that all the women in our sample are pregnant at the first observation and we then use the fact that the mothers were asked annually about any experience of abuse for the first seven years of the survey. There are pros and cons to using data on mothers with young kids.

One drawback is that most learning about partners can be expected to take place early on in relationships, but many of the women in our sample will have lived with their partners for several years before entering the survey. Nevertheless, about half of all the women in the sample have been living with their current partners for no more than three years at the beginning of the survey. Also on the positive side, the years following the birth of a child is a key period when women’s decisions regarding further fertility and if and when to return to work are particularly

salient.

A key consideration in the recent literature on domestic violence has been the measurement of abuse with several authors advocating strict objective measures (Aizer, 2010; Tertilt and van den Berg, 2015). The emphasis on having an objective measure is natural in contexts where the research aim is to understand the effect of various factors on the incidence of abuse, as findings could otherwise be confounded by reporting and other composition effects.

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 will make use of a relatively subjective 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. Though these questions are much less specific than used in many dedicated domestic violence survey modules, we will nevertheless show that estimated incidence of abuse in our sample is very similar – both in terms of level and in terms of demographic correlates – to the best available evidence from the UK drawn from the Crime Survey for England and Wales.

Our model is exclusively focused on the behaviour of women. The behaviour of their male partners with respect to abuse is modelled in a highly “reduced form” by assuming exogenous probabilities of engaging in abuse depending on his nature and on the woman’s chosen action, notably her labour supply. Our model thus assumes that men’s behaviour is non-strategic. This modeling choice is done in part for simplicity, but also in part as response to lack of consensus in the literature regarding the drivers behind male abuse.

Indeed, some researchers suggest that male abuse may be either due lack of self-control (Card and Dahl, 2011), or it may represent rational behaviour (Aizer, 2010); some have suggested that males use abuse “instrumentally” to extract resources from the victims family (Bloch and Rao, 2002), the affect bargaining power (Ramos, 2016) or to induce preferred actions (Anderberg and Rainer, 2013), while others researchers assume that violent actions represent direct preferences for abuse (Bowlus and Seitz, 2006).

We model women’s choice of partnership status, labour supply, and child-bearing from the

moment they enter the “marriage market” until the end of their fertile period. As such, our model builds on an established literature developing lifecycle models of family decisions (van der Klaauw, 1996; Francesconi, 2002; Keane and Wolpin, 2010; Gemici and Laufer, 2014). The relationship between our work and two contributions to this literature are worth noting in more detail. The first is Brian, Lillard, and Stern (2006). Their key focus is on the choice between marriage and cohabitation, and they treat labour supply and fertility as exogenous. In their model, a given partnership is associated with an underlying true match quality – a continuously distributed variable – and, in each period, the couple (but not the researcher) observe and enjoy an unbiased signal of that match quality. Thus a couple learn their match quality over time. The learning setting in our model is on the one hand simpler: women learn their partner’s type with only a binary type space, and, importantly, belief updating is based on the experience of abuse which is observable in the data. On the other hand, by endogenizing fertility and labour supply we study key behavioural responses to learning beyond partnership decisions.

The second is Bowlus and Seitz (2006) which is the only contribution to date that estimates a lifecycle model with domestic violence. In their model, men rationally decide if and when to be abusive based on their preferences for violence. However, as the authors assume that women always know their partner’s abuse preferences there is no learning. Moreover, they take fertility as exogenously given. Their model is then estimated on cross-sectional data with retrospective self-reported information. Our model departs from their work by modeling learning, endogenous fertility, estimating on actual dynamic data, and by not imposing an assumed rationality on male abusive behaviour.

The paper is outlined as follows. Section II describes the ALSPAC sample and we present a set of linear fixed-effects regressions to highlight some key dynamic relationships in the data. Section III we describe the model, starting with a simple illustrative version highlighting in particular how key parameters will be identified from the onset rate, persistence and overall level of abuse before outlining the full empirical model.

Section IV outlines the estimation approach while Section V reports the parameter estimates and the model fit. Section VI present several counterfactual experiments, exploring the impact

on women’s behaviour and experience of abuse of, *inter alia*, women’s preferences and information, and of changes in the economic environment, including the removal (e.g. through stricter sentencing policy) of serial abusers, higher female wages, and more generous child support enjoyed by single mothers. Section VII concludes.

II Data and Illustrative Dynamics

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.¹ The initial recruits were pregnant women with estimated dates of delivery between April 1991 and December 1992.² While first and foremost a child development survey, ALSPAC also repeatedly surveyed the mothers of the study children (and their partners), and it is from these mothers’ surveys that our data is constructed. Our female sample population is hence the mothers of the ALSPAC children who were repeatedly surveyed as part of the ALSPAC design. In particular we will use that the mothers were surveyed roughly annually about key events in their lives, including their experience with abuse, up until when the survey child was about 6 years old, a maximum of seven observation years for each female respondent.³ A particular feature of our data is hence that, per construction, each woman in the sample has a birth between the first and the second wave. Ethical approval for the study was obtained from the ALSPAC Ethics and Law Committee and the Local Research Ethics Committees.

ALSPAC recruited 14,541 pregnant women who returned at least one questionnaire or attended a “Children in Focus” clinic by 19/07/99.⁴ In order to conduct our analysis, we impose

¹For a detailed description of the ALSPAC cohort, see Boyd et al. (2013).

²Ethical approval for the study was obtained from the ALSPAC Ethics and Law Committee and the Local Research Ethics Committees.

³The 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 key questions when the study child was aged 8, 21, 33, 47, 61 and 73 months respectively. After that the key abuse-related question were no longer regularly asked.

⁴Of these initial pregnancies, there was a total of 14,676 fetuses, resulting in 14,062 live births and 13,988

a set of restrictions on the sample. Avon is the South-West region of England where, in the UK context, the population is known to be of predominantly white ethnic origin (ONS, 2012). In order to avoid issues with small cell sizes, we drop all women who are of Asian, Black or other/unknown ethnic origin, dropping a further 2,614 women. We then remove all women for whom basic demographic information on age and/or academic qualification level is missing, dropping a further 508 women.

If a respondent misses one or more survey within our observation window, we only retain person-year observations up until the first missing survey, and we then only keep the women observed in pregnancy and for at least one wave post birth, dropping 672 women. We then eliminate person-year observations with missing information on the key time-varying variables: partnerships status, births, and abuse which eliminates a further 1,312 women. We further eliminate women who were pregnant with the ALSPAC child below the age of 17 (32 women) or above 40 (44 women) in order to be consistent with our lifecycle model below. This leaves a sample of 9,359 women, with a total of 56,926 person-year observations, with over 80 percent of the sample women observed for the complete seven years.

Sample Population

We start by characterizing the demographic characteristics of the sample population at baseline.⁵ Note that at this stage, all the women in the sample are mid-pregnancy. Panel A of Table 1 gives basic information about the population at this stage. The sample women were, on average, 28 years old at the start of the survey..

The vast majority, 96 percent, of the women lived with a male partner at baseline, and had done so for over four and a half years on average. 55 percent of the sample women already had

children who were alive at 1 year of age. When the oldest children were approximately 7 years of age, an attempt was made to bolster the initial sample with eligible cases who had failed to join the study originally. As our study only makes use of data up until the age of 6, we do not use these additional study women – known as the “Phase I enrolment” sample – in our analysis.

⁵The ALSPAC study website contains details of all the data that is available through a fully searchable data dictionary. See <http://www.bristol.ac.uk/alspac/researchers/access/>

Panel A: Sample Population at Baseline					
Variable	Mean	Std. Dev.	Variable	Mean	Std. Dev.
Age in Years	28.11	4.55	Nr Children	0.782	0.895
Has Partner (“Married”)	0.962	0.192	Low Qualification	0.244	0.430
Years with Partner	4.84	3.53	Medium Qualification	0.381	0.486
Any Child	0.553	0.497	High Qualification	0.374	0.486
Obs.	9,359				
Panel B: Female Wages by Age and Qualification					
Age Group	Mean	Std. Dev.	Qualification	Mean	Std. Dev.
Aged 17-24	5.55	1.79	Low Qualification	5.37	1.62
Aged 25-31	6.47	2.38	Medium Qualification	6.05	1.89
Aged 32-45	7.49	2.87	High Qualification	8.46	2.88
Obs.	56,790				
Panel C: Male Wages by Age and Qualification					
Age Group	Mean	Std. Dev.	Qualification	Mean	Std. Dev.
Aged 17-24	7.10	2.19	Low Qualification	7.14	2.25
Aged 25-31	8.62	3.12	Medium Qualification	8.67	3.15
Aged 32-65	9.88	3.58	High Qualification	9.94	3.63
Obs.	53,326				

Table 1: Demographic characteristics of the ALSPAC sample.

at least one child at baseline and the average number of existing children was 0.78.

Primarily for estimation purposes, we delineate only a limited number of qualification groups of roughly equal size. The “low” qualification group include women without any formal qualification or with a qualification at NVQ1 level, most notably a CSE or a “low” GCSE.⁶ The “medium” qualification group hold a qualification at NVQ2 level, most notably an O-level degree or “high” GCSE. The “high” qualified group hold a qualification at NVQ3 level or beyond, which includes A-level degree, university undergraduate degree and beyond.

⁶The sample population were potentially affected by two major UK educational reforms. First, the raising of the school leaving age from 15 to 16 in 1973 affecting those born after September 1957. This reform is well-known to have significantly raised the academic qualification rate (Dixson and Smith, 2011). Second, the introduction of the General Certificate for Secondary Education (GCSE) in 1986, affecting those born after September 1970. This reform, which replaced the previous age 16 qualifications known as the Certificate for Secondary Education (CSE) and the the General Certificate of Education Ordinary Level (O-level), further increased the academic qualification rate. In our sample, only about ten percent of the women were born early enough to face the lower school leaving age, and also only about ten percent of the women were born late enough to face the new GCSE system. Hence the sample women overwhelmingly faced a school-leaving age of 16 with the CSE/O-level qualification system. The O-level qualification in particular acted a pathway to the A-level (Advanced Level) and the A-level qualification in turn is the standard requirement for university entrance.




Figure 1: Distribution of age, partnership duration and number of children at baseline

Figure 1 provides further details of age, partnership duration and children at baseline. The left hand figure shows that many of the women were in their mid- to late 20s when entering the survey. The middle figure shows that 40 percent of the women in the sample had a current partnership duration of no more than 3 years. The right hand figure shows that about 45 percent of the women in the sample, the ALSPAC child represented a first birth, and a further 38 percent had only one previous child.

While the ALSPAC data unfortunately only contains information about total household income (including benefit income), it does contain detailed occupational information in the form of the standard SOC90 classification system at the 3-digit level. We use this information to impute an hourly wage for each person-year observation, based on the respondent’s most recent occupation in the listing of over 300 possible occupations. Specifically, for each occupation in the classification system, we compute and use the average wage among all women aged 18-59 observed in the UK Labour Force Survey between 1993 and 1999. Panel B in Table 1 provides summary statistics on these imputed wages by age and qualification. The wages of male partners are imputed in the same way using the partner’s occupation, and summary statistics by age and qualification is provided in panel C of Table 1.

The ALSPAC further only contains hours of work in banded form and we use this information to assign each observation to one of the three labour supply states – not working, working part-time, and working full-time – as outlined above. The model estimated below will focus on annual earnings. For that purpose we will assume that part-time and full-time work corresponds to 20

and 40 hours/week for 50 weeks/year, thus imputing annual earnings to be 1,000 and 2,000 times the hourly wage respectively.

Partnership Status, Children, Abuse and Labour Supply

For partnership status we make no distinction between marriage or cohabitation and refer to a woman “married” if she currently lives with a male partner either as married or cohabitating, and as “single” otherwise. We will correspondingly refer to the event of a woman leaving her partner as a “divorcing” and the event of forming a new partnership as “marrying”.

The vast majority of observed partners are also the natural father to the child that the woman is pregnant with at the start of the survey; however, we make no formal distinction between natural fathers and other male partners. In a small number of cases, a woman is observed to switch partner from period to the next.⁷ For estimation purposes we want to avoid direct partner-to-partner transitions; in such cases we therefore ignore the initial months of the new partnership and effectively assume that the woman was single for one intervening period. Panel A of Table 2 notes that, across all person-year observations, only some 94 percent of women are married. This is obviously lower than at baseline; indeed, the married rate drops monotonically over time and reaches 90 percent by the end of the sample period. Panel A further notes that the overall divorce rate is little less than 2 percent, whereas single women marry at an annual rate of 12 percent.

The birth dummy variable indicates the event of a birth between the previous and the current period. All women in the sample, per construction, give birth to the ALSPAC child between the first and the second period. The birth rates reported in Panel B of Table 2 are therefore computed using data from period three onwards. As such it measures the arrival of subsequent siblings to the ALSPAC child. Nearly half of the women in the sample have some further birth in the years that follow and the average birth rate from sample period 3 onwards is 0.12. The table shows that a woman is less likely have a birth in any given period if she had one in the

⁷Direct partner to partner transitions can be detected in the data from information provided by the mother on the duration of her current relationship and on the status of the male partner being the biological father of the ALSPAC child or not.

Panel A: Partnership Status				
Time t	Mean	Time $t + 1$		
		Single	Married	
Single	0.063	0.880	0.120	
Married	0.937	0.019	0.981	
Obs.		56,926		
Panel B: Birth Incidence				
Time t	Mean	Time $t + 1$		
		No Birth	Birth	
No Birth	0.879	0.856	0.144	
Birth	0.121	0.926	0.074	
Obs.		37,876		
Panel C: Labour Supply				
Time t	Mean	Time $t + 1$		
		Not Working	Working PT	Working FT
Not Working	0.471	0.801	0.166	0.033
Working PT	0.345	0.183	0.703	0.114
Working FT	0.184	0.229	0.302	0.469
Obs.		53,746		
Panel D: Abuse Incidence				
Time t	Mean	Time $t + 1$		
		Not Abused	Abused (any)	
Not Abused	0.908	0.943	0.057	
Abused (any)	0.092	0.505	0.495	
Time t	Mean	Time $t + 1$		
		Not Physically Abused	Physically Abused	
Not Physically Abused	0.976	0.982	0.018	
Physically Abused	0.024	0.647	0.353	
Time t	Mean	Time $t + 1$		
		Not Emotionally Abused	Emotionally Abused	
Not Emotionally Abused	0.913	0.945	0.055	
Emotionally Abused	0.087	0.511	0.489	
Obs.		56,926		

Table 2: Summary statistics and transition rates for variables measured over time.

previous period, reflecting that the spacing of births is typically more than one year. Children born within the sample period are added to each woman's existing children at baseline, thereby keeping track of how many children she has at any moment in time.⁸

Information on hours of paid work is available in each wave and we use this information

⁸The focus on own biological children to the female respondent thus means that we include children who potentially have left home and but not any children of the partner who may reside with the household. These issues are likely to be relatively minor. First, since each woman is pregnant at the beginning of the sample period, few of them will have children old enough to have moved out. Second, as a stylized fact, the vast majority of children from separated parents live with their natural mothers.

to classify the female participant’s current labour supply status as not-working, working part-time or working full-time, where the latter two categories are defined as working 1 – 25 or 25 hours/week or above, respectively. Part-time work is common in the data, across all periods. Full time work on the contrary has a stronger time profile. About 40 percent of the women work full time at baseline. This then drops sharply in conjunction with the birth of the ALSPAC child before gradually picking up again over time. By the end of the sample period, close to a quarter of the women are in full time paid work. This feature of the data should also be kept in mind when interpreting the observed transition rates in panel C of Table 2. Notably, the fact that the majority of women observed in full time employment have left this state by the following period is a reflection of them reducing labour supply in conjunction with a birth.⁹ Also, the low rate of direct transitions from being out of the labour force to full time employment reflects that many of the women in the sample re-enter employment more gradually via part time employment.

As noted above, our indicators of abuse are based on self-reported measures. At each wave the mother was asked to complete a 42-item recent-events inventory.¹⁰ Two recurrent items were “Your partner was physically cruel to you” and “Your partner was emotionally cruel to you” and we take the responses at face value. For the majority of the analysis we will combine the two into a single indicator of abuse of any kind, but panel D of Table 2 presents a breakdown also by type of abuse. Overall 9.2 percent of women report some form of abuse in any give year, with nearly all those reporting some abuse also reporting emotional abuse. The fraction of women reporting physical abuse is significantly lower at 2.4 percent. A striking feature of the abuse variables is their persistence: half of those reporting some abuse in a given period also report abuse in the following period.

Figure 2 shows how the reported incidence of abuse varies by age group and qualification level. As is well-known, younger women tend to be more at risk than older women, and less qualified women are more at risk than the more qualified.

⁹Conditioning on no child birth between t and $t + 1$, the rate of remaining in full time employment is 83 percent.

¹⁰Each questionnaire specifies to the respondent what time period is meant by “recent”; in particular these periods are specified so as to measure events since the last survey.

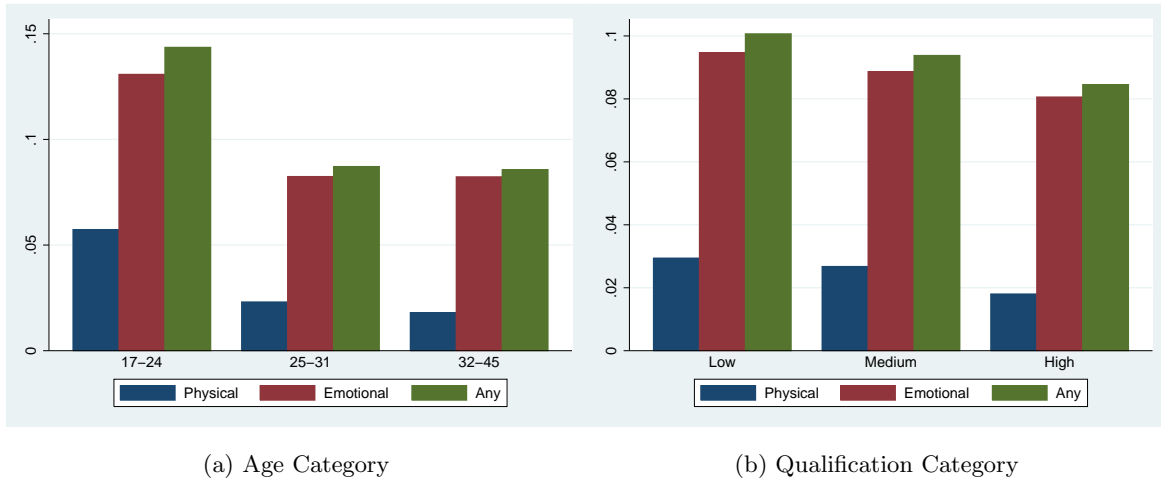


Figure 2: Incidence of abuse by demographic group.

Even though the measures we use are self-reported and subjective we show in Appedix A that, both in terms of level and demographic pattern, they agree well with the best available measures of physical and emotional abuse obtained from the interpersonal violence modules of the Crime Survey for England and Wales.

Illustrative Dynamics

In order to guide our modelling of women’s responses to abuse, we will start with a preliminary analysis of the dynamic patterns in the data. Noting however that all women in the sample, per construction, report a birth between the first and the second sample period, the below illustrative analysis will be entirely based on person-year observations from the third sample period onwards when the ALSPAC child would have been aged between 20 months and 7 years. As noted above, this is a time when many of the women in the sample made key choices in terms of either returning to work or having a further child, and also a period when a number of them chose to break up their current partnerships. We will use a set of simple linear regressions – estimated both by pooled OLS and with individual fixed effects – to explore the association between these choices and the incidence of abuse. In doing so it is important to pay attention to the timing of the variables involved as some variables – most notably marital and labour supply status – measure the state of a variable *at a given point in time*, whereas other variables

– including abuse, births and divorce – measure events occurring *over the 12 months*.

For ease of interpretation all models are estimated as simple linear probability models. All models estimated by OLS include dummies for qualification level, and all regressions include controls for the female respondent’s age and age squared. The results are presented in Table 3.

Consider first how current marital status at time t relates to the experience of abuse. Since the abuse reported at t indicates events over the past 12 months, we can relate the respondent’s current marital status to her currently reported abuse experience. However, for comparison, we further include abuse reported at $t - 1$ (thus measuring exposure to abuse 13-24 months prior to the currently observed marital status). The first columns of Panel A of Table 3 reports the results from a simple linear OLS regression whereas the second column gives the results from a corresponding individual fixed effects (within) regression. Both regressions indicate that a woman is markedly more likely to be single at time t if she also report having experienced abuse at some point between time $t - 1$ and t or between $t - 1$ and $t - 2$. The lower estimated coefficients in the FE model suggests potential selection both into partnerships and partnership responses to abuse.

In order to focus on the choice of separating from a partner as a response to abuse, the remaining columns in Panel A use only observations the respondent was married at $t - 1$ and we use as dependent variable whether she divorced her partner between $t - 1$ and t . In order to ensure that we only relate this to abuse that predates the potential divorce decision, we only include lagged abuse, that is abuse occurring between $t - 2$ and $t - 1$. Hence the regression considers whether, among all women who were married at $t - 1$, those who were abused between $t - 2$ and $t - 1$ were more more likely to subsequently divorce between $t - 1$ and t . Columns (iii) and (iv) report the results from an OLS and FE regression respectively, with both indicating a positive effect of abuse on divorce risk. The final two columns in Panel A look separately at physical and emotional abuse, again estimated with fixed effects. Both indicate a positive impact on divorce risk, though the impact of lagged physical abuse is imprecisely measured.

The FE regression in specification (iv) suggests a clear divorce response to abuse: using the estimated coefficients, the model predicts that the divorce hazard increases from 1.8 percent to

Panel A: Partnership Status						
Dep. Var.	Married at t		Divorced since $t - 1$			
Specification	(i)	(ii)	(iii)	(iv)	(v)	(vi)
Any Abuse (t)	-0.097** (0.008)	-0.065** (0.008)				
Any Abuse ($t - 1$)	-0.133** (0.008)	-0.062** (0.007)	0.063** (0.005)	0.030** (0.006)		
Physical Abuse ($t - 1$)					0.022 (0.012)	
Emotional Abuse ($t - 1$)						0.031** (0.006)
Obs.	36,641	36,641	34,482	34,482	34,482	34,482
Method	OLS	FE	OLS	FE	FE	FE
Panel B: Birth						
Dep. Var.	Birth since $t - 1$					
Specification	(i)	(ii)	(iii)	(iv)	(v)	(vi)
Any Abuse ($t - 1$)	-0.046** (0.005)	-0.027** (0.007)				
Physical Abuse ($t - 1$)			-0.035** (0.010)	-0.011 (0.012)		
Emotional Abuse ($t - 1$)					-0.047** (0.005)	-0.027** (0.007)
Obs.	35,033	35,033	35,033	35,033	35,033	35,033
Method	OLS	FE	OLS	FE	OLS	FE
Panel C: Labour Supply Status						
Dep. Var.	Not Working at t		Working PT at t		Working FT at t	
Specification	(i)	(ii)	(iii)	(iv)	(v)	(vi)
Any Abuse (t)	-0.005 (0.008)	-0.018 (0.010)	-0.013 (0.008)	0.015 (0.010)	0.018** (0.006)	0.003 (0.007)
Obs.	31,485	31,485	31,485	31,485	31,485	31,485
Method	OLS	FE	OLS	FE	OLS	FE
Panel D: Abuse						
Dep. Var.	Any Abuse since $t - 1$		Physical since $t - 1$		Emotional since $t - 1$	
Specification	(i)	(ii)	(iii)	(iv)	(v)	(vi)
Working PT ($t - 1$)	-0.007 (0.003)	0.009* (0.004)	-0.001 (0.002)	0.007** (0.003)	-0.006 (0.003)	0.008 (0.004)
Working FT ($t - 1$)	0.033** (0.006)	0.026** (0.007)	0.009** (0.003)	0.011** (0.004)	0.032** (0.005)	0.026** (0.007)
Nr Children ($t - 1$)	0.014** (0.002)	0.027** (0.005)	0.007** (0.001)	0.009** (0.003)	0.014** (0.002)	0.026** (0.005)
Obs.	33,015	33,015	33,015	33,015	33,015	33,015
Method	OLS	FE	OLS	FE	OLS	FE

Table 3: Illustrations of the dynamic pattern in the data using pooled OLS and fixed-effects regressions.

4.8 percent. The fact that the regression focuses only on *non-immediate* separation responses to abuse – that is, it does not account for abuse followed by a separation *within* the same time period – implies that this is, on the one hand, almost certainly an underestimate of the divorce response to abuse. On the other hand, the rate of divorce between $t - 1$ and t among women who also report abuse over that same period is about 13 percent (not in table), and is almost certainly an overestimate of the divorce response to abuse. The data therefore clearly indicates that the vast majority of women who experience abuse do not, at least in the short-run, leave their partners.

Consider next how the experience of abuse affects the decision to have a (further) child. Since the birth variable indicates a birth event over the last year we lag the abuse variables by one period. The regressions reported in Panel B thus relate a birth occurring between $t - 1$ and t to whether the woman experienced abuse between $t - 2$ and $t - 1$. The regressions in this panel further controls for the existing number of children at $t - 1$. Recalling that the average probability of a further birth in the periods included in the regressions is 0.12 (see Table 1), the first two columns suggest that an experience of abuse reduces the fertility hazard by 20 - 40 percent. The final four columns report negative coefficients both for physical and emotional abuse, though the coefficient on the former is small and not very precisely estimated in the FE specification. A consistent pattern is again that the estimated effects of abuse are smaller in the FE specifications than in the pooled OLS specification, suggesting selection effects based on unobserved heterogeneity.

Panel C looks at how a woman's labour supply status at time t is affected by the experience of abuse between $t - 1$ and t . The regressions in this panel further control for the lagged labour supply status and current number of children. The evidence here is rather mixed, but if anything the results suggest that women respond to experiencing abuse by less frequently remaining out of the labour force. The results in panels A-C thus suggest that women who experience abuse respond by more frequently leaving their partner, reducing their fertility, and possibly also increasing their labour supply. However, while we have here explored each of these choice dimensions separately, the overall response may well involve a combination of responses,

which will be accounted for in the structural model estimated below.

In panel D, we switch focus to the determinants of abuse, in particular labour supply and kids. Since the abuse variable indicates the event of abuse over the past year, we lag the explanatory variables so as to measure the impact of the state at $t - 1$ on abuse experience between $t - 1$ and t . The potential impact of labour supply on exposure to abuse in particular is a widely discussed topic in the literature.¹¹ The regressions presented here strongly suggest that female labour supply, particularly in the form of full-time employment, increases exposure to abuse: all regressions, whether estimated as pooled OLS or with individual fixed-effects, and for any or either type of abuse, suggest that working full time is associated with about 20-30 percent higher risk of abuse than when not working. The results for part-time work is less conclusive, with the estimated signs differing between the OLS and FE. When estimated with fixed-effects, the results suggest that working part-time may also be associated with more exposure to abuse, but the estimated coefficients always smaller in absolute terms than for full time employment.

The regressions further show a positive association between number of children and abuse risk. Indeed, the estimated coefficient is generally larger in the FE regressions than in the OLS. To interpret this, it should be kept in mind that the effect of children in the FE model is identified from within-respondent variation. Hence the coefficient in the FE model effectively captures the effect of a recent birth on abuse. This is thus different from the OLS estimated effect which is identified from the cross-sectional variation in the female respondents' number of kids.

III 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. We assume that there are two types of males: (i) men who have a “violent nature” and who are abusive with a high frequency, and (ii) men who have a “non-violent nature” and who are abusive much more rarely. The behaviour of males is taken as exogenously given and random. While men differ in their nature,

¹¹See e.g. Macmillan and Gartner (1999) for a seminal contribution and Heise (2011) for a recent review.

a woman who meets a new prospective partner does not directly observe his nature; instead she forms beliefs which she updates based on her observations of his behaviour. In particular, when experiencing abuse, her belief that he has a violent nature increases, which in turn lowers her expected future utility from remaining in her current partnership.

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 if she only discovers that the husband has a violent nature after having invested in children. This in turn implies a strategic incentive to defer fertility within a new relationship in order to first become confident about the partner’s nature.

We also allow for the possibility that the rate at which men with a violent nature in particular are abusive varies with the woman’s level of labour supply. E.g. a woman who experiences abuse – and thus increases her belief that the partner has a violent nature – may choose to increase her labour supply in order to build up her work experience and hence future earnings capacity, anticipating that she is now more likely to leave her partner. The incentives for doing so while still married will critically depend on whether increasing her labour supply will increase or decrease the risk of further abuse.

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 structure allows us to replicate key features in the data relating to the incidence of abuse.

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, $m_t \in \{0, 1\}$. In this simple version we normalize the utility of being single to zero and let ψ^m denote the per-period utility of being married. In addition, each woman obtains, in each period t , a random utility ε_t^m from being married which we take to be i.i.d. normally distributed with zero mean and variance σ_m^2 .

A woman who enters a period as married can choose to either remain married or to divorce. Single women randomly receive marriage offers at rate ζ from new prospective partner. Any

new prospective male partner is of one of two possible types, $r \in \{0, 1\}$: he is either of the “non-violent type” ($r = 1$) or he is of the “violent type” ($r = 0$). The husband’s type is a fixed personal characteristic. However, the woman receiving the marriage offer does not observe the proposing male’s type. The probability that a new prospective male partner is of the non-violent type is denoted by

$$\phi_b = E[r] \in (0, 1), \quad (1)$$

and thus also represents the woman’s initial beliefs about the type of any new partner.

What distinguishes male types is their propensity to engage in abuse. Let $z_t \in \{0, 1\}$ indicate a woman’s exposure to abuse at time t and let χ_r denote the per period probability that a male of type r engages in abuse; we then assume that $0 < \chi_1 < \chi_0 < 1$. This difference in abuse behaviour means that a woman updates her beliefs based on the husband’s observed actions. 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

$$\phi_{t|z_{t-1}=0} = \frac{\phi_{t-1}(1 - \chi_1)}{\phi_{t-1}(1 - \chi_1) + (1 - \phi_{t-1})(1 - \chi_0)}, \quad (2)$$

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

$$\phi_{t|z_{t-1}=1} = \frac{\phi_{t-1}\chi_1}{\phi_{t-1}\chi_1 + (1 - \phi_{t-1})\chi_0}. \quad (3)$$

Experiencing abuse is associated with the 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

$$\pi(\phi_t) = \phi_t\chi_1 + (1 - \phi_t)\chi_0, \quad (4)$$

captures her perceived likelihood of experiencing abuse.

Consider then a woman who is either married or who has met a new potential partner. Based on her current beliefs, $\phi_t \in [0, 1]$, about her available partner and also on her marriage utility shock ε_t^m she decides on her marital status, $m_t \in \{0, 1\}$, for the current period. Letting δ denote the discount rate, the model can then be solved using standard 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.¹²

¹²Formally, $V^m(\phi_t)$ and V^s satisfy

Consider then divorce behaviour. A woman who enters a period as married with beliefs ϕ_t will divorce if

$$\psi^m + \varepsilon_t^m - \pi(\phi_t)\psi^z + \delta \left[\pi(\phi_t)V^m(\phi_{t+1|z_t=1}) + (1 - \pi(\phi_t))V^m(\phi_{t+1|z_t=0}) \right] < \delta V_s, \quad (5)$$

which means that there will be a threshold ε_t^m below which she will divorce. Moreover, this threshold value will be a function of her current beliefs ϕ_t : women with more pessimistic beliefs about their husbands' types will set a higher threshold value for ε_t^m and will hence be more prone to divorce. As experiencing abuse will worsen a woman's beliefs about her husbands nature, divorce will be more likely after incidents of abuse.

For a given set of parameters, the model can be solved numerically and then forward-simulated to generate a steady state distribution of marital status, abuse incidence and beliefs. Doing so allows us to highlight how key parameters of the model relate to moments in the data. In this simple model we normalize ψ^m to unity and set the discount parameter to $\delta = 0.95$. A woman's rate of accepting a new marriage offer will be the same as the rate at which a married woman with belief ϕ_b continues her marriage. Since this rate can be expected to be high (see below), the rate at which single women enter new partnerships is largely determined by the partner meeting rate ς . E.g. at $\varsigma = 0.14$, the expected duration of singlehood will, empirically plausibly, be around seven years.

We now turn to the more specific parameters, σ_m^2 , ψ^z , ϕ_b , χ_0 and χ_1 , and discuss how these can be related to divorce behaviour and abuse incidence.

The regressions in Table 3 showed that women who reported experiencing abuse were more likely to divorce: the raw divorce rates in the data are 0.075 and 0.014 (see Table 6 below) which we match here. These stylized fact help pin down the variance σ_m^2 and the disutility ψ^z . As the utility shock ε_t^m is temporary, it needs to be sufficiently large to make even some women who hold very positive beliefs about their husbands occasionally choose to divorce. Moreover, even women who were exposed to abuse are distinctly more likely to remain with their partners than divorce which effectively limits ψ^z . Setting $\psi^z = 0.32$ and $\sigma_m^2 = 2.72$ generates steady

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

and $V^s = \varsigma V^m(\phi_b) + \delta(1 - \varsigma)V^s$ respectively.

state divorce rates that matches the empirical moments. More generally it implies that the probability of divorce as a function of beliefs ϕ_t goes from around a low rate of little over one percent for women who firmly believe that their partners are of the non-violent type up to close to ten percent for women firmly believe that their partners are of the violent type.

This suggests that the systematic instantaneous utility of marriage is substantially reduced by abuse, but remains positive, $\psi^m - \psi^z > 0$. Women who, through experience, firmly believe that their partners are of the violent type do not necessarily immediately leave their husbands; however, compared to women who hold more positive beliefs, the abused women are less willing to accept temporary negative marriage utility shocks.

The three remaining parameters are the baseline beliefs (or, equivalently, the frequency of non-violent males among prospective partners), and the abuse rates of non-violent and violent male types respectively. These are closely related to the overall abuse rate and the abuse “transition” rates in Panel D of Table 2. Setting $\phi_b = 0.64$, $\chi_1 = 0.03$ and $\chi_0 = 0.71$ generates an overall abuse rate of 0.092, an abuse onset rate of 0.057 and a persistence rate of 0.495.

In order to generate level a persistence of abuse corresponding to that observed in the data it must be some men are high-repeat offenders. However, it cannot be the all abuse is perpetrated by such violent men. In particular, the higher divorce rate after abuse implies that the prevalence of violent men in the steady state pool of husbands is much lower – less than 10 percent in the current example – than the prevalence of such men among the new potential partners encountered by single women ($1 - \phi_b = 0.36$). In order to still predict a substantial overall rate of abuse, the calibrated model suggests a low rate of abuse also by men with non-violent nature.

The empirical model presented below will expand on the current one by incorporating also labour supply and fertility decision. As such it will have a different cardinalization and will have a finite time horizon. Nevertheless, it will retain some of the key qualitative features from this very simple framework. In particular, it will have a similar size of the temporary marriage utility shocks relative to the systematic utility of marriage, a similar disutility of abuse relative to the baseline utility of marriage, and similar estimated frequencies of male types and their




Figure 3: The Evolution of the Belief Distribution

abuse propensities.

One key feature that the full model below will share with the current simple model is the basic properties of the learning environment. Indeed, as will be seen below, when we estimate the full model we will obtain that the share of non-violent types among proposing men is of a similar magnitude as found here. Likewise, the full model will find a similar rates of abuse as here (though slightly lower abuse towards working women). In particular, this means that the “speed” of learning and the associated distribution of beliefs in the full estimated model is very similar to that generated in the current simple model.

Given that, we will use the current model to illustrate the nature of the distribution of beliefs and speed of learning. Consider then a population of women who get married at time $t = 0$ and who, for simplicity, remain married for at least three periods. Figure 3 then illustrate how the distribution of beliefs develop in this population. In the figure, a blue line represents a period without abuse and a red line a period with abuse. At the time of marriage all women hold the baseline beliefs. In the first period of marriage 73 percent of the women experience no abuse and update their belief to 0.86; in contrast, 27 percent of women do experience abuse and update their beliefs to 0.07. After two periods of marriage, each woman will either have experienced 0 (63 percent), 1 (19 percent) or 2 (18 percent) episodes of abuse, leading to beliefs that are 0.95,

0.20 and 0.003 respectively. After one further period, 88 percent of the women have beliefs that are either above 0.98 or below 0.02. Hence within just a two or three periods, the vast majority of women will have very firm beliefs about the nature of their partners. This high speed of learning also means women can expect not to have to wait long if for instance they choose to delay fertility until they have are more confident about their partner's nature.

The Full Empirical Model

The full version of the model that we take to the data models women's choices with respect to marital status, employment status and child-bearing in a finite horizon setting, focusing on choices made between the ages of 17 and 41. In each period t 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 is 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 maximizes her present value of lifetime utility, discounted at rate δ . The utility flow in period t is specified as

$$U_t = \frac{\mu^{k_t} C_t^{1-\lambda}}{1-\lambda} + (\Psi_t^m - \bar{\Psi}_t^z) m_t + \Psi_t^n, \quad (6)$$

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 normalized 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 women 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 captures the direct utility of children, Ψ_t^n . The Ψ -terms will be further specified below.

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

$$C_t = \begin{cases} \tau (w_t + w_t^h - c_t) & \text{if } m_t = 1 \\ w_t - c_t & \text{if } m_t = 0 \end{cases}, \quad (7)$$

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 childcare costs incurred at t (specified

further below).

Wage Offers

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

$$w_t^k = \exp\left(\beta_0^k + \beta_1^k a + \beta_2^k x_t - \beta_3^k x_t^2 + \varepsilon_t^k\right), \text{ for } k = 1, 2, \quad (8)$$

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 earnings factor a is assumed to be stochastically related to her observed educational attainment level, which, as described in Section II, is either "low", "medium", or "high", $q \in \{0, 1, 2\}$. We specify the relationship between q and a to be logistic,

$$\frac{\Pr(a = 1|q)}{\Pr(a = 0|q)} = \exp(\beta_0^a + \beta_1^a d_{q=1} + \beta_2^a d_{q=2}), \quad (9)$$

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

$$x_{t+1} = x_t + k_t, \quad (10)$$

starting 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 , $k = 1, 2$.

The husband's earnings in (7) is specified in a similar way as

$$w_t^h = \exp\left(\beta_0^h + \beta_1^h a + \beta_2^h x_t^h + \varepsilon_t^h\right), \quad (11)$$

where ε_t^h is also also a temporary productivity shock. The presence of the woman's permanent productivity type a in the husband's wage offer equation (11) captures a systematic spousal wage correlation, representing assortative mating on ability. Married couple tend also to be similar in age (note) and since men are assumed to always be working FT in our model, the partner's experience x_t^h increases lineary with the woman's age (beyond 17).

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

$$A = \begin{bmatrix} a_{11} & 0 & 0 \\ a_{21} & a_{22} & 0 \\ a_{h1} & 0 & a_{hh} \end{bmatrix}. \quad (12)$$

Childcare costs are assumed to depend non-linearly on the number of children, n_t , on the mother's level of labour supply, and on her marital status. Specifically, c_t is specified as

$$c_t = \rho^{k_t} (\beta_1^c n_t + \beta_2^c n_t^2) - \beta_3^c n_t (1 - m_t), \quad (13)$$

where ρ^{k_t} indicates what proportion of these costs are incurred at labour supply level k_t . We normalize so that $\rho^2 = 1$ for full time employment, and constrain ρ^0 and ρ^1 to be in the unit interval. β_3^c is a reduction in the costs per child for single mothers. As this represents income obtained by a single mother independently of her chosen labour supply, a natural interpretation of β_3^c is that it represents a flat child support payment per child from the father.

Marriage, Learning and Conception

The marriage and learning side of the model is exactly as in 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 $\varsigma \in (0, 1)$, with men being of two possible types, $r \in \{0, 1\}$.¹³ The fraction of encountered men who are of the non-violent type is ϕ_b , and $z_t \in \{0, 1\}$ indicates whether or not the woman is exposed to abuse in period t . z_t is realized after the woman has decided on her level of labour supply k_t and conception f_t . Hence a married woman makes these decisions under uncertainty about potential exposure to abuse. A non-violent husband type is, in any given period, abusive with probability $\chi_1 \in (0, 1)$ while a violent husband type is abusive with probability $\chi_0^{k_t} \in (0, 1)$ where the superscript k_t indicates that we now allow the probability of abuse to vary with the woman's chosen level of labour

¹³Note that we are not using any time subscript on the husband's type to indicate that his type is fixed. Nevertheless, it should be clear that if a woman remarries, her next husband may be of a different type.

supply. A woman's beliefs are updated exactly as in (2) and (3) while also taking into account that the abuse rate by violent men depends on her chosen labour supply.

The expected disutility from abuse for a married woman in (6) with current belief ϕ_t and chosen labour supply k_t is given by $\bar{\Psi}_t^z = \pi(\phi_t, k_t) \psi^z$ where

$$\pi(\phi_t, k_t) = \phi_t \chi_1 + (1 - \phi_t) \chi_0^{k_t}, \quad (14)$$

is her perceived probability of experiencing abuse and where ψ^z is the direct disutility of 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

$$n_{t+1} = n_t + f_t. \quad (15)$$

The direct utility from children and conception in (6) is specified as

$$\Psi_t^n = \beta_1^n n_t - \beta_1^n n_t^2 + f_t \varepsilon_t^f, \quad (16)$$

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

$$\Psi_t^m = \psi^m + \varepsilon_t^m, \quad (17)$$

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 of each other.

IV Estimation

The model is estimated using method of simulated moments (McFadden, 1989; Pakes and Polard, 1989). This approach entails, for any trial parameters, first solving the model using backwards induction and then forward-simulating to obtain simulated panel data with lifecycle paths for a large number of individuals with a distribution of observable characteristics that correspond to that observed in the data.

Simulated Population and Sampling

We model choices of women between the ages of 16 and 44, and generate – for any trial parameters – simulated outcomes for 15,000 women with a distribution of academic qualifications – the only source of observed initial heterogeneity – as observed in the data. We then focus on the simulated outcomes between the ages 17 to 40 to help correct for the initial conditions problem and end-of-horizon effects. We further limit this subset to women who have had at least one child, calculating the simulated moments amongst these women starting from the year in which they give birth. In order to closely replicate the sampling scheme in the ALSPAC data, we focus on women who have at least one child, calculating the simulated moments amongst these women starting from the year in which they first give birth. Hence, childless women do not contribute to the calculation of simulated moments.

Our strategy produces only two negligible deviations from the ALSPAC sampling scheme. First, the ALSPAC’s initial observation for each woman is during pregnancy, while we “sample” women from the time they give birth. In the model, conception choice and birth occur within the same period (year). Conception is at the beginning of the period and birth at the end. Therefore, we cannot separate the timing of these two events in the model. Second, women entering the ALSPAC sample are not necessarily carrying their first child.

Standard Errors

Standard errors are obtained by taking the square root of the diagonal elements

of the variance-covariance matrix $Q_S(W)$,

$$Q_S(W) = \left(1 + \frac{1}{S}\right) \left[\frac{\partial b(\theta_0)'}{\partial \theta} W^* \frac{\partial b(\theta_0)'}{\partial \theta} \right]^{-1}, \quad (18)$$

where $\partial b(\theta_0)' / \partial \theta$ is the first derivative of the vector of moments b with respect to the parameter vector θ . S is the number of simulations (15,000 * 24) and W is the weighting matrix. We use the identity matrix for W and set $1/S = 0$, given the large number of simulations ($1/S = 0.000003$). Use of the identity matrix rather than an ideal weighting matrix only reduces efficiency. $\partial b(\theta_0)' / \partial \theta$ is numerically approximated using parameter bump sizes that vary between .01% and 1% depending on the sensitivity of the moments.

Identification

Overall 41 parameters are estimated using 85 empirical moments. The set of moments included in the estimation, which contain both static and dynamic ones and ones that link choice dimensions, can be broadly split into three main groups by what they help identify. The first group contains moments related to employment (by age and qualification, and transitions) and wages (by labour supply status and qualification, and of husbands). These moments strongly identify the parameters associated with the wage offer functions, unobserved ability structure, the disutility of work effort, income associated with non-employment, and the correlation between per-period earnings shocks.

The second group of moments correspond to those used for the simple illustrative model above and thus contain moments related to marriage and abuse: marriage rates and marital transitions rates, the overall incidence and transitions in exposure to abuse, the divorce response to abuse, and abuse by labour supply status. These moments identify the disutility of abuse, the size of match quality shocks, the arrival rate of partners, and the type-specific abuse frequencies. Combined with the identified earnings structure, the observed marriage rate further identifies the sharing parameter. Interestingly, observed rates of abuse help identify the marital utility shock, which has been difficult to identify in discrete choice dynamic programming (DCDP) models that do not incorporate domestic abuse data (see, e.g., Keane and Wolpin (2010) and Sauer (2015)).

The third main group of empirical moments relates to children and contains fertility measures (completed and timing), the fertility response to abuse, the rate of out of wedlock births, and employment status by marital status. These moments help identify the utility of children, conception utility shocks, child-related costs, and the level of child-support.¹⁴

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

¹⁴As 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 the Office for National Statistics (REF).

parameter of relative risk aversion λ is set at 0.7. Identification of δ and λ is a common problem in DCDP models.¹⁵

V Estimation Results

In this section, we present estimates of the structural model presented in Section III. In assessing the model, we consider the within-sample fit and the reasonableness of the parameter values.

Moments and Model Fit

Tables 4 to 6 present the moments included in the estimation, comparing the empirical and simulated values. Table 4 presents on the employment-related moments – labour supply status by age, qualification level and marital status, and employment transitions. Table 5 presents hourly wages by labour supply status, of husbands and by qualification level. Table 6 presents the moments related to marital status, fertility and abuse incidence.

The model fits the labour supply pattern quite well, though the predicted age-gradients of labour supply are somewhat larger than observed in the data. For the labour supply by marital status, the model captures the lower rate of employment by single mothers and the relatively high frequency of part-time work by married mothers. In terms of employment transitions, of those who enter employment, the model somewhat overpredicts the rate at which women enter full-time employment. Also, among those leaving part-time employment, the model slightly overpredicts the rate of moving up to full-time employment rather than leaving employment.

In terms of hourly wages, the model correctly predicts that the accepted wages of full time workers exceed those of part time workers, though it over- and under-predicts the standard deviation of male and female earnings respectively.

¹⁵A woman's beliefs are a key state variable and the belief space, $\Phi \in [0, 1]$, is in principle continuous. For computational purposes we use a 74-point grid. The grid used is not equi-spaced, but rather denser towards the ends of the unit interval. This feature is chosen to reflect the natural properties of the learning process in which belief changes implied by the Bayesian updating process tend to be smaller when the prior is closer to either zero or unity.

Panel A: Employment Status			
	Not Working	Working Part-Time	Working Full-Time
All	0.471 <i>0.516</i>	0.345 <i>0.327</i>	0.184 <i>0.157</i>
Panel B: Employment Transitions			
	Not Working at $t + 1$	Part-Time at $t + 1$	Full-Time at $t + 1$
Not Working at t	0.801 <i>0.846</i>	0.166 <i>0.121</i>	0.033 <i>0.034</i>
Part-Time at t	0.183 <i>0.160</i>	0.703 <i>0.736</i>	0.114 <i>0.105</i>
Full-time at t	0.229 <i>0.198</i>	0.302 <i>0.316</i>	0.469 <i>0.485</i>
Panel C: Employment Status by Age Group			
	Not Working	Working Part-Time	Working Full-Time
Aged 17-24	0.585 <i>0.769</i>	0.207 <i>0.110</i>	0.208 <i>0.121</i>
Aged 25-31	0.486 <i>0.561</i>	0.344 <i>0.287</i>	0.170 <i>0.153</i>
Aged 32-40	0.438 <i>0.411</i>	0.374 <i>0.419</i>	0.188 <i>0.170</i>
Panel D: Employment Status by Marital Status			
	Not Working	Working Part-Time	Working Full-Time
Single	0.590 <i>0.467</i>	0.240 <i>0.328</i>	0.171 <i>0.205</i>
Married	0.463 <i>0.521</i>	0.352 <i>0.327</i>	0.185 <i>0.152</i>
Panel E: Employment Status by Qualification Level			
	Not Working	Working Part-Time	Working Full-Time
Low Qual.	0.575 <i>0.622</i>	0.307 <i>0.289</i>	0.118 <i>0.089</i>
Medium Qual.	0.490 <i>0.586</i>	0.349 <i>0.300</i>	0.160 <i>0.115</i>
High Qual.	0.396 <i>0.316</i>	0.362 <i>0.403</i>	0.242 <i>0.281</i>

Table 4: Matched moments: employment

Panel A: Accepted Hourly Wages by Labour Supply Status and of Husbands			
	Part-Time	Full-Time	Husband
Mean	6.86 <i>6.81</i>	7.90 <i>8.13</i>	9.40 <i>9.38</i>
St. Dev	2.70 <i>2.20</i>	2.90 <i>2.99</i>	3.51 <i>3.32</i>

Panel B: Accepted Hourly Wages by Qualification Level			
	Low Qual.	Medium Qual.	High Qual.
Mean	5.35 <i>6.49</i>	6.07 <i>6.85</i>	8.78 <i>7.96</i>
St. Dev	1.64 <i>2.44</i>	1.92 <i>2.54</i>	2.89 <i>2.45</i>

Table 5: Matched moments: wages

The model fits the marital transitions well, though the overall divorce rate is slightly under-predicted. The model slightly under-predicts the age at first birth, but predicts an age pattern in out-of-wedlock births. The empirical annual birth rates presented in panel C are for the periods following the birth of the ALPAC child and hence capture births of younger siblings. The simulated moment is computed in the corresponding way as births of further children after the birth that triggers inclusion in the data used for computing the simulated moments. The model also predicts well the proportion of women who remain childless and the distribution of number of children among those who have children.

Turning to abuse, in line with the simple model above, the full empirical model replicates quite closely the overall level of abuse the onset and persistence of abuse. The model predicts a realistic gap in abuse incidence between low and high qualified women. It also predicts that the exposure to abuse falls with age, though the model struggles to replicate the very high incidence in the youngest age group. It is important to remember that the age- and qualification gradients of abuse are generated directly from the structure of the model, without any assumptions for instance about differential tolerance of abuse.

Importantly, the model replicates the U-shaped relationship between the level of labour supply and exposure to abuse, implying that PT work is the labour supply status least associated with abuse. Finally, the model predicts well the relative divorce and fertility choices of abuse

Panel A: Marriage Rate and Marital Transitions			
	Mean	Single at $t + 1$	Married at $t + 1$
Single at t	0.063 <i>0.098</i>	0.880 <i>0.875</i>	0.120 <i>0.125</i>
Married at t	0.937 <i>0.902</i>	0.019 <i>0.015</i>	0.981 <i>0.985</i>

Panel B: Age at First Birth and Proportion Giving Birth as Single		
Average Age at 1st Birth	Proportion of Out of Wedlock Births	
	Age 17-24	Age 25-40
26.9	0.123	0.028
<i>25.6</i>	<i>0.106</i>	<i>0.066</i>

Panel C: Fertility			
Birth Rate Mean	Birth Rate Married	Birth Rate Single	Average Nr Children
0.121 <i>0.121</i>	0.126 <i>0.127</i>	0.037 <i>0.065</i>	1.82 <i>1.85</i>

Panel D: Distribution of Nr Children			
Childless	1 Child	2 Children	3+ Children
0.190 <i>0.172</i>	0.102 <i>0.140</i>	0.409 <i>0.349</i>	0.299 <i>0.339</i>

Panel E: Abuse Rate and Abuse Transitions			
	Mean	No Abuse at $t + 1$	Abuse at $t + 1$
No abuse at t	0.908 <i>0.913</i>	0.943 <i>0.955</i>	0.057 <i>0.045</i>
Abuse at t	0.092 <i>0.087</i>	0.505 <i>0.471</i>	0.495 <i>0.529</i>

Panel F: Abuse Rate By Qualification Level		
Low Qual.	Medium Qual.	High Qual.
0.101 <i>0.104</i>	0.094 <i>0.092</i>	0.085 <i>0.063</i>

Panel G: Abuse Rate By Age Group		
Age 17-24	Age 25-32	Age 33-40
0.144 <i>0.095</i>	0.087 <i>0.082</i>	0.085 <i>0.088</i>

Panel H: Abuse Rate By Labour Supply at $t - 1$		
Not Working	Part-Time	Full-Time
0.101 <i>0.098</i>	0.084 <i>0.057</i>	0.106 <i>0.111</i>

Panel I: Choice at t by Abuse Experience at $t - 1$			
Divorce Rate		Birth Rate	
Abuse at $t - 1$	No Abuse at $t - 1$	Abuse at $t - 1$	No Abuse at $t - 1$
0.075 <i>0.044</i>	0.014 <i>0.012</i>	0.075 <i>0.054</i>	0.126 <i>0.128</i>

Table 6: Matched moments: marriage, fertility and abuse.

and non-abused women.

Parameter Estimates

The estimated parameters are presented in Tables 7 and 8, with Table 7 presenting the β -coefficients from equations (8), (11), (16), (13) and (9) and Table 8 presents all remaining parameters. Before discussing the estimated values, we will present the empirical moments used in the estimation and the estimated model's fit to these moments.

Consider now the estimated parameters, starting with those presented in Table 7. The rate of earnings growth with respect to experience ranges from over 8 percent in the early career states down towards zero for women who have worked full time for most of their adult lives. The coefficient on age in the male earnings equation gives the estimated annual earnings growth of males.

The estimated coefficients on ability indicate that high ability women can earn nearly twice the amount of low ability women when both are working full time, and about 60 percent more when working part time. The difference in their husbands' earnings are however proportionately smaller, with the husbands to high ability women earnings about 25 percent more than the husbands to low ability women. The estimated child-related costs are substantial, ranging from about £3,000 per year with one child to over £7,000 with three children. The child support payment per child is only about a third of the child-related costs for a first child.

The probabilities of being low ability ($a = 0$) if unqualified, medium qualified, and high qualified are 0.38, 0.21 and 0.19 respectively. Hence the low ability women are a minority group concentrated among, about half of whom are drawn from the group of unqualified women.

Consider now the parameters presented in Table 8. Similar to the simple model above, the estimated ψ^z indicates that exposure to abuse substantially reduces the utility from marriage, and the estimated σ_m^2 indicates that there are sizeable temporary match quality shocks. The estimated prevalence of non-violent types ϕ_b is also very similar to the simple model above, and the abuse rates by non-violent and violent men are also similar to the simple model, with some variation by the woman's labour supply for the violent men. The estimated meeting rate ζ is also effectively unchanged from the simple model. Childcare costs are nearly eliminated for women

Panel A: Wage Offer Functions				
	Non-Emp. $\log(w^0)$	PT Emp. $\log(w_t^1)$	FT Emp. $\log(w_t^2)$	Husband $\log(w_t^h)$
Constant	7.307 (0.00)	7.236 (0.00)	7.593 (0.00)	9.520 (0.00)
a		0.377 (0.00)	0.583 (0.00)	0.043 (0.00)
x_t		0.105 (0.000)	0.105 (0.000)	
$x_t^2/100$		-0.171 (0.0??)	-0.171 (0.0??)	
age_t				0.017 (0.00?)
$age_t^2/100$				-0.013 (0.00?)

Panel B: Child-Utility and Child-Related Costs		
	Utility Ψ_t^n	Cost c_t
n_t	0.715 (0.00)	4,735.29 (49.47)
n_t^2	-0.063 (0.00)	-234.07 (17.00)
$n_t(1 - m_t)$		2,899.08 (28.42)
$n_t^2(1 - m_t)$		-603.88 (28.42)

Panel C: Ability Probability Function	
Constant	0.255 (0.210)
$q = 1$	0.678 (0.133)
$q = 2$	0.999 (0.067)

Table 7: Parameter estimates: linear equations

Panel A: Preference Parameters					
Marriage		Abuse	Fertility	Work Effort Cost	
ψ^m	σ_m^2	ψ^z	σ_f^2	μ_1	μ_2
301.13	705.18	141.14	2.049	0.999	0.971
(1.52)	(9.45)	(2.81)	(0.023)	(0.001)	(0.001)
Panel B: Abuse Parameters					
Type Prob		Abuse Prob			
ϕ_b	χ_0	χ_1^0	χ_1^1	χ_1^2	
0.663	0.019	0.718	0.566	0.560	
(0.028)	(0.000)	(0.000)	(0.001)	(0.001)	
Panel C: Sharing, Cost Fractions, Meeting Rate					
Sharing	Childcare		Meeting Pr.		
τ	ρ^0	ρ^1	ς		
0.566	0.023	0.357	0.130		
(0.026)	(0.006)	(0.007)	(0.003)		
Panel D: Cholesky Terms					
a_{22}	a_{32}	a_{33}	a_{h2}	a_{hh}	
0.190	0.230	0.049	0.333	0.032	
(0.005)	(0.014)	(0.011)	(0.067)	(0.098)	

Table 8: Parameter estimates continued: remaining parameters

who do not work and substantially lower for women who work part time compared to women who work full time. The “sharing” parameter τ indicates close to equal sharing.¹⁶

VI Counterfactual Experiments

In this section we report on a number of counterfactual simulations. The first three simulations involve altering the financial incentives facing women in the model. In the final simulations we then alter the information available to women when they meet new potential partners. We start by describing how the environment is altered in each simulation.

Simulation 1: Eliminating the Gender Wage Gap

In this simulation 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

¹⁶It should be noted however that τ can also capture household public goods whereby the sum of her consumption as proportion of total household income (τ) and his corresponding consumption as a proportion of total income can exceed unity.

representing a gender-specific wage subsidy.

Simulation 2: A Marriage Bonus

Here we provide a direct bonus to married couples in the form of a lumpsum addition to the household budget. The level of the bonus was set at 50 percent of the average earnings of employed women in the baseline model (£4,746). This experiment could be taken to represent a policy that either directly provide tax relief to married couples or mitigates marriage penalties in the tax-benefit code.

Simulation 3: More Generous Child Support

Here we increase the child support parameter, β_3^c , by 20 percent relative to the baseline.

In the final two simulations we provide the women in the model with information about the men that they meet. To parameterize this, suppose that when a woman meets a new prospective partner she receives a signal, $s \in \{0, 1\}$, which is potentially correlated with the male’s true type. She can then decide whether or not to marry the male in question based on her beliefs after observing the signal. To model the precision of the signal, let

$$\Pr(s = 1|r = 1) = \frac{1 + \epsilon}{2} \text{ and } \Pr(s = 1|r = 0) = \frac{1 - \epsilon}{2}, \quad (19)$$

with parameter $\epsilon \in [0, 1]$. The baseline model corresponds to $\epsilon = 0$ whereby s is completely uninformative. At the other extreme, $\epsilon = 1$, whereby the signal satisfies $s = r$ for certain and hence is completely revealing the male’s type. For values of ϵ strictly in the unit interval s and r are correlated with ϵ parameterizing the informativeness of the signal. The last two columns in present the result from simulations where $\epsilon = 1$ (“complete information”) and $\epsilon = 0.5$ (“incomplete information”). The former case in particular completely shuts down learning since any woman will be fully informed about any male’s types from the very moment they meet.

Simulations 4-5: Perfect Information ($\epsilon = 1$) and Imperfect Information ($\epsilon = 0.5$)

Whereas in the estimation we focused on the population of mothers in order to match to the ALSPAC data, we will report simulated outcomes for the entire population.

Table 9 give the main results for the full population of women. Figure 4 highlights some more details of the dynamics behind these results, in each case contrasting four key simulations against the baseline economy.

	Baseline Model	Perfect Information	Increased Female Earnings	Marriage Bonus	Increased Child Support
Age at First Marriage	22.169	22.255	22.173	22.158	22.171
Proportion Married	0.689	0.685	0.687	0.692	0.688
Divorce Rate	0.021	0.021	0.022	0.020	0.021
Divorce Rate after Abuse	0.054	0.055	0.055	0.051	0.054
Age at First Birth	25.584	23.282	26.280	23.726	22.402
Proportion Childless	0.172	0.169	0.150	0.150	0.062
Average Nr of Children: All	1.854	2.117	1.886	2.007	2.485
Low Qualified	2.124	2.334	2.107	2.350	2.758
Medium Qualified	2.015	2.248	2.009	2.199	2.639
High Qualified	1.489	1.822	1.594	1.556	2.124
Non-Employed	0.344	0.323	0.249	0.433	0.540
Working Part-Time	0.175	0.239	0.220	0.173	0.148
Working Full-Time	0.481	0.438	0.531	0.395	0.312
Average Own Earnings	6,573	6,968	8,645	5,729	4,936
Average Husb. Earnings	18,499	18,514	18,499	18,499	18,501
Abuse Frequency: All	0.117	0.101	0.114	0.121	0.120
Low Qualified	0.120	0.104	0.117	0.125	0.125
Medium Qualified	0.117	0.103	0.114	0.121	0.121
High Qualified	0.114	0.097	0.112	0.118	0.115

Table 9: Counterfactual Simulations: Lifetime Outcomes

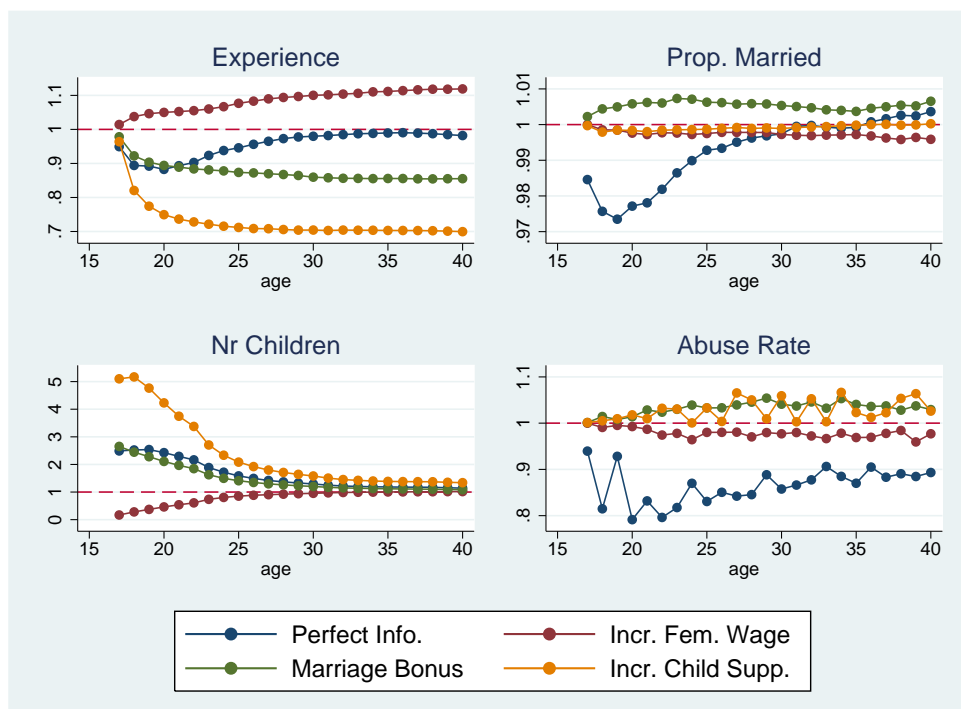


Figure 4: Parameter estimates continued: remaining parameters



Figure 5: Counterfactual Simulations: Labour Supply by Qualification Group

Figure 5 give further details for how labour supply vary across qualification groups.

We will discuss each simulation in turn.

Eliminating the Gender Wage Gap

Relative to the baseline model, higher female wages encourages women to work more as seen in Figure 4. Labour market experience grows faster than in the baseline economy, and by age 40 average labour market experience is more than 10 percent higher than in the baseline economy. Figure 5 further shows that the increase in labour supply comes in particular from an increased labour force participation among the low- and medium qualified women.

It also delays fertility. From Table 9 we also see that the average age at first birth increases by 0.7 years and Figure 4 shows that below the age of 25, the average number of children is well-below that in the baseline economy. Over time, fertility catches up so that so that, by age 40, slightly fewer women remain childless and the average number of children is also slightly higher. Improved female wages however does not delay entry into marriage as that is mainly determined by the random availability of well-matched partners. Indeed, the average age at

first marriage is unaffected (Table 9) as is, by and large, the proportion of women married at any age (Figure 4). The delay in fertility, but not in marriage, implies that fewer births occur out-of-wedlock.

Turning to abuse, we see that the overall incidence of abuse decreases by a small amount – about 0.3 percentage points. This decrease is driven by the increased labour force participation. The result that higher wages for women would decrease exposure to abuse is in line with the findings of Aizer [complete]

A Marriage Bonus

A marriage bonus creates incentives to be married. However, in line with the findings in the literature, marriage behaviour is not particularly responsive to financial incentives (REF). As a result, age at first marriage, the probability of being married, and the probability of divorce change only marginally. However, as women generally expect to be married throughout most of their adult lives, a marriage bonus creates an expected income effect. This in turn reduces women’s focus away from their labour market careers and towards children. As can be seen from Table 9 and Figure 5 this is particularly pronounced in the lower two qualification groups.

These responses are visible in Figure 4. In particular, the figure shows that, by age 40, women have accumulated about 14 percent less labour market experience than in the baseline economy. They are also more likely to be married at any given age. Women have a higher average number of children at any given age, but most markedly so below the age of 30. Indeed, from Table 9 we see that the average age at first birth is reduced by close to two years, the fraction of women who remain childless decreases somewhat, and the average number of children by age 40 increases by close to 10 percent. As women advance their fertility more than their marrying, the proportion of children born out of wedlock increases.

The average rate of abuse increases by about 0.4 percent, mainly as a result of women more frequently being out of the labour force, but also in part due to them being less prone to divorce after experiencing abuse.

More Generous Child Support

At first glance, more generous child support could potentially enable mothers to "avoid" bad marriages, and as such it could potentially help reduce their exposure to partner violence. However, marriages can also be "bad" in terms of match quality, which would tend increase the attractiveness of single-motherhood in general. Moreover, more generous child support allows a single mother to boost her income, thereby boosting fertility in general – and out-of-wedlock fertility in particular – and also lowering labour supply.

These broader responses are visible in Figure 4. This figure shows that, with more generous child support, women accumulate less experience, have more children overall and particularly at an early age, and they are slightly less likely to be married. In Table 9 we see that the age at first birth reduces substantially, the average number of children (by age 40) increases by more than 0.6, and the fraction of women who remains childless drops sharply. As before, the age at first marriage is, in contrast, largely unaffected. As in the previous two scenarios, the impact on labour supply are particularly pronounced in the among women with low- or medium qualifications (Figure 5)

Turning to abuse, we see that there is an overall increase of 0.3 percent, affecting mainly women with low- and medium qualifications, reflecting their reduced labour force participation.

Information

In the final experiment we provide women with information about the nature of any prospective new partner that they meet. In the extreme case, the information is perfect – meaning that the male's type is immediately observable – and we will focus on this case here. As will be seen below, the case of a noisy signal generates outcomes that is in most case between that of no information and perfect information.

There are two immediate behavioural differences when women have perfect information relative to the baseline case. With upfront information about the nature of any new prospective partner, women can now in principle reject marriage offers from violent males. As a consequence, as highlighted in Figure 4, the rate at which women are married drops below age 30 and age

at first marriage increases somewhat as seen in Table 9.¹⁷ Over time, however, the proportion married catches up as divorces later in marriage in particular decreases.

An equally marked difference in behaviour lies with fertility: when women do not have to gradually learn about their partners' nature they delay fertility less. Indeed, indeed a majority of women who marry a non-violent man conceive a child immediately in the first year of marriage. In contrast, the fertility rate among women who accept marriage offers from violent men is very low. The net effect is to increase early fertility on average, leading to a timing pattern for fertility that is very similar to the case with a marriage bonus. Indeed, Table 9 highlights how the average age at first birth decreases by little over two years – about the same time it takes most women to become confident about the nature of their partners in the baseline model. As before, a decrease in the age at first birth tends to also increase the proportion of out of wedlock births. However, compared to the scenarios with a marriage bonus or an increased child support – where the age at first birth drop in a similar fashion – in this case the increase in the proportion of out of wedlock births increases relatively less, reflecting that fertility now takes place earlier in marriage.

With women having more children early on, they also supply less labour in the early stages of their careers. However, in contrast to the case of the marriage bonus where there was a permanent income effect, in the case of perfect information total accumulated labour market experience catches up with the baseline economy over time.

Table 9 reveals that, when full information becomes available, the largest increase in fertility occurs among the high qualified women. Correspondingly, in terms of labour supply, all qualification groups reduce their full-time labour supply and increase their part-time labour supply, where the latter effect is particularly pronounced among the medium qualified women.

Marriages to violent males not only have low fertility by also tend to be short-lived due to a significantly higher divorce rate. As a result, the proportion of women married to partners

¹⁷The drop in married rates is rather modest. This reflects that, even with perfect information, many marriage offers from violent men are accepted. This may seem paradoxical, but note that the rejection rate of marriage offers must, per construction, be of a similar magnitude to the divorce rate which is quite low even when women are certain or near certain that the partner has a violent nature.

with a violent nature decreases. As a result, the abuse rate reduces by about 14 percent relative to the baseline model, and the reduction in abuse is fairly uniform across qualification groups. Note also that this reduction is not due to labour supply changes as there is little change in the rate of non-employment.

The final column of Table 9 gives the results for an imprecise signal. In most cases, the results are inbetween the case of no information and the case of full information.

VII Conclusions

Women who experience abuse from their intimate partners almost certainly did not expect that when they first got together. In this paper we have estimated a learning model where women update their beliefs about their partner's nature during the course of the relationship. To this end we have used longitudinal data where women subjectively report their experience of any abuse, physical or emotional and where we observe partnership choices, fertility choices and labour supply choices.

The longitudinal nature of the data reveals that abuse – if it occurs – has a high degree of persistency. This is consistent with some men being high-repeat offenders, which is a central finding of the estimated model. But the model also shows that women respond to abuse by more frequently leaving their partners and by choosing not to have more children. But women's learning about the nature of their partners does not only lead to reactive behaviour. It also provides incentives for strategic behaviour. For instance, it creates incentives for delaying fertility, taking time to first ascertain the nature of the partner.

After estimating the model we illustrate these various effects – and various implied policy effects – through counterfactual simulations where we alter the economic environment, for instance by assuming that women can immediately observe their partner's nature.

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.

References

- Aizer, A. (2010), 'The gender wage gap and domestic violence', *American Economic Review* **100**, 1847–1859.
- Aizer, A. & Dal Bó, P. (2009), 'Love, hate and murder: Commitment devices in violent relationships', *Journal of Public Economics* **93**(3), 412–428.
- Alesina, A., Brioschi, B. & Ferrara, E. L. (2016), 'Violence against women: A cross-cultural analysis for Africa', NBER Working Paper No. 21901.
- Anderberg, D. & Rainer, H. (2013), 'Economic abuse: A theory of intrahousehold sabotage', *Journal of Public Economics* **97**(C), 282–295.
- Anderberg, D., Rainer, H., Wadsworth, J. & Wilson, T. (2016), 'Unemployment and domestic violence: Theory and evidence', *Economic Journal* **126**(597), 1947–1979.
- Angelucci, M. (2008), 'Love on the rocks: Domestic violence and alcohol abuse in rural Mexico', *B E Journal of Economic Analysis & Policy* **8**(1).
- Bloch, F. & Rao, V. (2002), 'Terror as a bargaining instrument: A case study of dowry violence in rural India', *American Economic Review* **92**.
- Bobonis, G. J., González-Brenes, M. & Castro, R. (2013), 'Public transfers and domestic violence: The roles of private information and spousal control', *American Economic Journal: Economic Policy* **5**(1), 179–205.
- Bowen, E., Heron, J., Waylen, A., Wolke, D. & ALSPAC Team (2005), 'Domestic violence risk during and after pregnancy: findings from a British longitudinal study', *BJOG: An International Journal of Obstetrics & Gynaecology* **112**, 1083–1089.
- Bowlus, A. & Seitz, S. (2006), 'Domestic violence, employment, and divorce', *International Economic Review* **47**, 1113–1149.
- Boyd, A. et al. (2013), 'Cohort Profile: The "Children of the 90s" – the index offspring of the Avon Longitudinal Study of Parents and Children', *International Journal of Epidemiology* **42**, 111–127.

- Brien, M. J., Lillard, L. A. & Stern, S. (2006), 'Cohabitation, marriage, and divorce in a model of match quality', *International Economic Review* **47**(2), 451–494.
- Card, D. & Dahl, G. (2011), 'Family violence and football: The effect of unexpected emotional cues on violent behavior', *Quarterly Journal of Economics* **126**, 103–143.
- Clark, L. E., Allen, R. H., Goyal, V., Raker, C. & Gottlieb, A. S. (2014), 'Reproductive coercion and co-occurring intimate partner violence in obstetrics and gynecology patients', *American Journal of Obstetrics and Gynecology* **210**, e1–8.
- Cormier, M. (2009), 'Domestic violence occurrences and the timing of welfare payments', M.A. Thesis, University of Ottawa.
- Dixson, M. & Smith, S. (2011), 'What determines the return to education: an extra year or a hurdle cleared?', *Economics of Education Review* .
- Enander, V. & Holmberg, C. (2008), 'Why does she leave? The leaving process(es) of battered women', *Health Care for Women International* **29**, 200–226.
- Erten, B. & Keskin, P. (2016), 'For better or for worse? Education and the prevalence of domestic violence in Turkey', Mimeo, Wellesley College.
- Francesconi, M. (2002), 'A joint dynamic model of fertility and work of married women', *Journal of Labor Economics* **20**(2), 336–380.
- Gemici, A. & Laufer, S. (2014), 'Marriage and cohabitation'. Royal Holloway, University of London.
- Heise, L. (2011), 'What works to prevent partner violence? An evidence overview', Working Paper, STRIVE Research Consortium, London School of Hygiene and Tropical Medicine,.
- Iyengar, R. (2009), 'Does the certainty of arrest reduce domestic violence? Evidence from mandatory and recommended arrest laws', *Journal of Public Economics* **93**(1), 85–98.
- Jasinski, J. L. (2004), 'Pregnancy and domestic violence: A review of the literature', *Trauma, Violence, and Abuse* **10**(9), 961–990.
- Keane, M. P. & 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* **51**(3), 851–892.
- Macmillan, R. & 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* **61**.
- McFadden, D. (1989), 'A method of simulated moments for estimation of discrete response models without numerical integration', *Econometrica* **57**(5), 995–1026.
- ONS (2012), 'Ethnicity and national identity in England and Wales: 2011'. Office for National Statistics report.
- Pakes, A. & Pollard, D. (1989), 'Simulation and the asymptotics of optimization estimators', *Econometrica* **57**(5), 1027–1057.
- Ramos, A. (2016), 'Household decision making with violence: Implications for conditional cash transfer programs', Universitat Autònoma de Barcelona.
- Riger, S. & Staggs, S. (2004), 'Welfare reform, domestic violence, and employment', *Violence Against Women* **10**(9).
- Tertilt, M. & van den Berg, G. (2015), 'The association between own unemployment and violence victimizations among female youth', *Journal of Economics and Statistics* **235**(4-5), 499–513.
- Tolman, R. M. & Wang, H.-C. (2005), 'Domestic violence and women's employment: Fixed effects models of three waves of women's employment study data', *American Journal of Community Psychology* **36**(1-2), 147–158.
- Tur-Prats, A. (2016), 'Family types and intimate-partner violence: A historical perspective', Mimeo, Universitat Pompeu Fabra.
- van der Klaauw, W. (1996), 'Female labour supply and marital status decisions: A life-cycle model', *The Review of Economic Studies* **63**(2), 199–235.
- Woodhouse, J. & Dempsey, N. (2016), 'Domestic violence in England and Wales', House of Commons Library, Briefing Paper Number 6337.

Appendix A: Comparison with the British Crime Survey

In this appendix we provide a comparison between the measured incidence of domestic abuse in the ALSPAC data and that in the British Crime Survey. In 1996 the British Crime Survey introduced a dedicated computerised self-completion intimate partner violence (IPV) module. The IPV module was further developed and used again in 2001, and then annually from 2005 onwards. Due to the assured anonymity and privacy when completing the module, and the detailed set of questions, the BCS IPV modules are considered to be one of the best quality large-scale evidence on the incidence of domestic abuse internationally.

The 1996 survey, while overlapping in time with the ALSPAC data, suffers two main drawbacks. First, it does not measure emotional abuse but instead focuses on threat of harm. Second, for physical abuse, it only gives details of the “most recent occasion”. In contrast, the later surveys contained an itemized list of abusive behaviours (see below), including verbal abuse and non-physical controlling behaviour, where the respondent was asked about any incidence of each type of behaviour over the past 12 months. We focus here on the 2001 BCS IPV survey as it offers a sufficient degree of detail while still being close in time to later years of the ALSPAC data. Including also surveys from 2005 onwards would substantially increase observation numbers, but would involve using survey data obtained on average over a decade after the ALSPAC sample.¹⁸ Hence we compare our ALSPAC sample to all women aged 17-45 in BCS 2001. This of course creates a key difference in that many of the women in the BCS sample are neither mothers nor pregnant. For this reason, we will present some comparisons that focus on the subsample of BCS women who have at least one child (“mothers”).

As part of the BCS IPV module, the respondents are asked if they have experienced any of the abusive behaviours listed in Table A.1 by an existing or past intimate partner over the past 12 months. We classify each recoded behaviour as either physical or non-physical abuse as indicated and create dummy variables to indicate the experience of one or more of the listed behaviours within each group. In addition to the IPV module questions, the BCS respondents are also queried about intimate partner abuse as part of the main BCS survey. The abuse

¹⁸An extended comparison that includes also BCS 2005 - 2007 is available on request from the authors.

Behavior	Physical Abuse	Non-Physical Abuse
Prevented from fair share of h-hold money		x
Stopped from seeing friends and relatives		x
Repeatedly belittled you		x
Frightened you, by threatening to hurt you		x
Pushed you, held you down or slapped you	x	
Kicked, bit, or hit you	x	
Choked or tried to strangle you	x	
Threatened you with a weapon	x	
Threatened to kill you	x	
Used a weapon against you	x	
Used other force against you	x	

Table A.1: Itemized abusive behaviours in the BCS IPV module.

reporting in the main survey is known to be substantially lower than in the dedicated IPV module, thus indicating under-reporting in the open survey. We include it here to compare the under-reported BCS measure to the ALSPAC measure.

Table A.2 provides summary statistics for the ALSPAC and the BCS data. The women in the data are similar in age and in the distribution of qualifications. However, all women in the ALSPAC data are either already mothers (or they are pregnant) whereby they naturally have a higher average number of children. The number of children become more similar when conditioning on having at least one child (“mothers”). As labour supply is strongly related to motherhood, the women in the ALSPAC data are significantly less likely to be working full time than the BCS women.

Turning to the measures of abuse, we see that average reported incidence of physical abuse is lower in the ALSPAC data than in the BCS. This is consistent with a degree of under-reporting. Nevertheless, the reported frequency of physical abuse in the ALSPAC is still noticeably higher than that in the BCS main survey, suggesting that the level of under-reporting in the ALSPAC is not severe.


In contrast, the reported rate of emotional abuse in the ALSPAC data is higher than the reported non-physical abuse in the BCS. This is consistent with the underlying questions in the ALSPAC survey being more open to interpretations by the respondent than the precisely

	ALSPAC	BCS 2001
Age	31.1 (4.98)	32.8 (7.41)
Qualification: Low	0.244 (0.430)	0.294 (0.456)
Qualification: Medium	0.381 (0.486)	0.301 (0.459)
Qualification: High	0.374 (0.486)	0.405 (0.491)
Nr of Children	1.85 (1.03)	1.15 (1.09)
Nr of Children (Mothers)	2.01 (0.92)	1.80 (0.83)
Not Working	0.471 (0.499)	0.319 (0.466)
Working PT	0.345 (0.475)	0.271 (0.445)
Working FT	0.184 (0.388)	0.410 (0.492)
Abuse Any	0.092 (0.289)	0.084 (0.277)
Physical Abuse	0.024 (0.153)	0.043 (0.203)
Emotional Abuse	0.087 (0.282)	0.062 (0.241)
Ph. Abuse: Main Survey		0.011 (0.103)
Nr. Obs.	56,926	2,142


Table A.2: Summary statistics for the ALSPAC sample and the BCS 2001 sample.

itemized questions in the BCS IPV module. Nevertheless, the physical and the emotional abuse variables are highly overlapping in both datasets: the correlation between physical and emotional abuse is 0.40 in ALSPAC and 0.38 in the BCS. Combining physical and emotional abuse into any abuse, gives a similar overall overall rate in the two data sets.


Figure A.1 highlights how the incidence of any abuse varies with the respondent's demographic characteristics in each data set. For comparability, we restrict the BCS sample to include only mothers. Panel (a) shows that the abuse incidence decreases with age in both data sets, while panel (b) shows that the rate of abuse is highest among low qualified women in both data sets. Finally, panel (c) highlights how there is a U-shaped relationship between labour supply and reported abuse in both data sets.



(a) Age group



(b) Qualification level



(c) Level of labour supply

Figure A.1: Incidence of any intimate partner abuse in the ALSPAC sample and in BCS 2001.