Women's Employment, Husbands' Economic Self-Interest and Domestic Violence

Deniz Sanin[†]

March 25, 2026

Abstract

This paper presents evidence that providing employment opportunities to women decreases domestic violence when the husband has economic self-interest in the wife's work capacity. I use the government-induced rapid expansion of the coffee mills in Rwanda in the 2000s, which increased the value of coffee farmer couples' output and provided wage employment for women. Since the mill operates only during the harvest months, the husband's cost of physically incapacitating his wife changes within the year. This variation, in conjunction with novel monthly administrative records on domestic violence hospitalizations, provides a way to distinguish the incapacitation cost mechanism from the rise in women's bargaining power and household income, which are identified by multiple rounds of household survey data. I replicate the findings using the universe of monthly domestic violence police records in Uganda over more than a decade, providing external validity for the incapacitation cost mechanism.

JEL Codes: J12, J16, J21, J31, O12
Keywords: Domestic violence, female labor supply, Rwanda

[†]Assistant Professor, Department of Economics, University of South Carolina. Email: deniz.sanin@moore.sc.edu.
This paper was previously circulated as "Paid Work for Women and Domestic Violence: Evidence from the Rwandan Coffee Mills". I am very grateful to Garance Genicot, Laurent Bouton, Mary Ann Bronson and Martin Ravallion for their invaluable guidance and support on this project. I sincerely thank Joe Altonji, Costas Meghir, Rohini Pande, and Mark Rosenzweig for their detailed comments and support during my visit at the Yale Economics Department and the Economic Growth Center, and Claudia Goldin and Larry Katz for their detailed feedback and support during my visit at the Harvard Economics Department. I owe special thanks to Andrew Zeitlin and Leodimir Mfura, who helped me in my data application. For very helpful comments, I would like to thank Daron Acemoglu, Marcella Alsan, Nava Ashraf, Sonia Bhalotra, Alberto Bisin, Lauren Bergquist, Kin Blackburn, Richard Blundell, Rosella Calvi, Raj Chetty, Jishnu Das, Kevin Donovan, Edward Glaeser, Penny Goldberg, Rema Hanna, Harvard Development and Labor Groups, Jonas Hjort, Reshma Hussam, Billy Jack, Supreet Kaur, Gabriel Kreindler, Eliana La Ferrara, Roger Lagunoff, Ben Marx, Ross Mattheis, Mushfiq Mobarak, Ameet Morjaria, Kaivan Munshi, Nathan Nunn, Orgul Ozturk, Alex Poirier, Nancy Qian, Dani Rodrik, Simone Schaner, Frank Schilbach, Daniel Valderrama, Alessandra Voena, Yale Development and Labor Groups, and participants at the NBER Development/BREAD Fall Meeting, NBER SI: Gender in the Economy, NEUDC and SOLE. This project has human subjects approval from Rwanda National Ethics Committee. I acknowledge the financial support of Georgetown University and the University of South Carolina. I also acknowledge Harvard Kennedy School Women and Public Policy Program as a supporting entity. All errors are my own.

1 Introduction

Domestic violence is an extreme form of gender inequality and a global health problem of epidemic proportions (WHO, 2013). 1 in every 3 women worldwide has experienced physical and/or sexual violence from their partners in their lifetime (WorldBank, 2015). Providing jobs to women is widely discussed as a policy to address the problem. Based on theories of domestic violence, employment can decrease violence by improving women's outside options and reducing financial stress through an increase in household income. Alternatively, it can increase violence due to husbands' incentives to extract women's resources or through male backlash. Existing empirical evidence supports both channels. However, these studies either focus on the effects of unearned income, such as cash transfers and dowry payments (Angelucci 2008, Bobonis et al. 2013, Hidrobo et al. 2016, Haushofer et al. 2019, Bloch and Rao 2002, Calvi and Keskar 2021), or they study employment primarily as an income shock due to their source of variation and data.[1] Yet, a job is more than income. It is the act of earning income by using one's capacity to work or (health) human capital (Becker 2007, Goldin 2016), and it cannot be performed if one is physically incapacitated—for instance, due to violence. If a husband's violence incapacitates his wife and creates an economic cost to him—because he benefits from her work—this cost itself can disincentivize him from engaging in violence.

This paper presents evidence that providing employment opportunities to women decreases domestic violence when the husband has economic self-interest in the wife's work capacity. I also provide evidence that an increase in women's outside options and their contribution to household income operate as additional mechanisms. To establish results, I use a natural experiment, the government-induced rapid expansion of the coffee mills in Rwanda in the 2000s, that affected coffee farmer couples. First, the opening of a mill enables the husband to sell the couple's harvest at a high premium on the international coffee market. This increases the value of the wife's physical capacity to perform the tasks she does as an unpaid family worker on the couple's plot. Second, a mill opening enables the wife to become a wage worker at the mill, giving her personal earnings, which she shares with her husband. Third, because the mill operates only during the harvest months, the husband's cost of physically incapacitating his wife changes within the year, given the economic incentives created by the first two points. This variation, in conjunction with monthly administrative records on domestic violence hospitalizations, provides a way to distinguish the

[1] These studies examine the effects of labor market shocks as a change in potential wages or the likelihood of having earnings (Aizer 2010, Anderberg et al. 2016, Erten and Keskin 2021, Erten and Keskin 2022). Bhalotra et al. 2021 is an exception, as the authors emphasize that their results on job-loss effects on domestic violence are consistent with both the loss of income and the increase in the time couples spend together. Kotsadam and Villanger 2022 conducts an RCT that randomizes job offers in Ethiopia. However, the experimental design focuses on identifying the effect rather than the underlying mechanisms or the non-monetary features of a job.

incapacitation cost mechanism from the rise in women's bargaining power and household income, which is identified by multiple rounds of household survey data.

In 2002, the Rwandan government adopted the National Coffee Strategy to shift to high-quality mill-processed coffee production to participate in the international specialty coffee market (Boudreaux, 2011). A public-private partnership project in the early 2000s helped farmers form cooperatives and build mills in their communities. Afterward, mill construction continued nationwide. Between 2002 and 2012, the number of mills increased from 5 to 213.[2]

A coffee mill processes coffee cherries, the harvested fruit of the coffee tree, into coffee beans for export. The context provides two key features for identification. First, the opening of a mill creates a spatial variation. A mill serves coffee farmers, mostly couples, who reside within its catchment area, approximately a 4 km radius circular buffer zone around the mill. This defined catchment area exists because coffee cherries begin to rot within a few hours of harvest if they are not transported to a mill in time. Second, a mill opening creates both yearly and monthly time variations. Before a mill opens, the wife harvests the coffee cherries with the husband and then processes them at home as a female-dominated task. The husband sells the home-processed coffee in the local market for a low price as a male-dominated task and keeps the income as his personal earnings. After a mill opening, the couple harvests and carries the cherries to the mill, and the husband sells them to the mill at a high premium.[3] The mill demands paid labor during the harvest months (March-July) for coffee processing tasks that its machinery cannot do. Since these tasks are female-dominated, a mill enables the wife in the catchment area to transition from an unpaid family worker in the couple's plot to a wage worker in the mill for the same tasks as before. As a result, the wife earns personal income, which she shares with her husband.[4]

Using a difference-in-differences (DID) event study design, I first identify the causal effect of a mill opening (the event of interest) by interacting spatial variation (individuals residing within and outside a mill's catchment area) with time variation (before and after a mill opening) for ever-married women and men (mostly their husbands). I capture time variation using the number of years an individual has been in the mill's catchment area. This strategy relies on self-reported domestic violence and labor market outcomes. To study the incapacitation channel, I then use novel monthly administrative records covering the universe of domestic violence hospitalizations. These data come from towards the end of the expansion period, when most mills had already been built

[2]In 2012, coffee accounted for nearly 20% of Rwanda's exports and 15% of its GDP (Macchiavello and Morjaria, 2020). As of 2018, there were 312 mills in the country.

[3]The husband earns a higher income from selling cherries to a mill compared to selling home-processed coffee in the local market (Macchiavello and Morjaria, 2020), since mill-processed coffee is of higher quality than home-processed coffee and is sold on the international coffee market. Farmers outside the catchment areas continue to process coffee at home after a mill opening, selling their harvest in the local market for a low price.

[4]In my data, 65% of women decide how to spend their earnings jointly with their husbands rather than independently, supporting the claim that women share their earnings with their husbands.

and the number of hospitals within a mill's catchment area was fixed. For this mechanism, I employ a different DID event study design, where the event of interest is the start of mill operations in March, marking the onset of the coffee harvest season. I exploit spatial variation (hospitals located within and outside a mill's catchment area) but now interact it with calendar months to capture the impact of mill operations within the year—restricted to the harvest months—on hospitalizations.[5] This introduces the second source of temporal variation in the paper.

The Rwandan Demographic Health Surveys (DHS) provide information on ever-married women and men's self-reported labor market outcomes and women's experiences of domestic violence in the past 12 months. The Hospital Management Information System (HMIS) contains the universe of monthly hospitalizations due to physical and sexual gender-based violence, where I use hospitalizations of women older than 18 as a proxy for domestic violence.[6] I combine data on the universe of mills from the Rwanda GeoPortal and Macchiavello and Morjaria (2020). Since all datasets are geocoded, I link individuals and hospitals to mills based on their GPS coordinates. Additionally, I use a household survey on individual earnings and household food consumption expenditure, along with DHS data on women's participation in household decisions, to study the mechanisms behind the results.

Being in the catchment area of a mill increases women's probability of working for cash in the past 12 months by 6 percentage points in both the 2–3 year and 4–6 year periods following the mill opening, relative to the year before, representing a 16% increase relative to the sample mean (0.37). The probability of working in the past 12 months remains unchanged for women. There is no change in aforementioned outcomes for men. To establish these results, I use individuals residing outside the catchment areas but within the same district as the mills as the control group. The results are robust to an alternative control group—individuals residing in the donut area between 4 and 8 km from the mills. Using additional household survey data, I also show that a mill opening increases women's and men's last daily earnings.

One year after a mill opening, being in the catchment area reduces women's probability of self-reporting physical or sexual domestic violence in the past 12 months, relative to the year before. The effect size declines over time and becomes statistically insignificant 4–6 years later. Examining physical and sexual domestic violence separately reveals a striking pattern. Physical domestic violence—including behaviors that could physically incapacitate women, such as kicking, punching, dragging, twisting an arm, or being threatened/attacked with a knife, gun or any other weapon—consistently decreases by 9 percentage points over time (0–1, 2–3, and 4–6 years later), a 29% reduction relative to the sample mean (0.31). In contrast, the effect on sexual domestic violence—includes behaviors relatively less likely to cause physical incapacitation—weakens

[5] Mill-exposed women plausibly seek care at hospitals within a mill's catchment area due to proximity.

[6] This is because 70% of women over 18 are married in Rwanda.

over time, becoming statistically insignificant after 4–6 years, contributing to the decline in the effect size of overall reported domestic violence.[7]

Is a woman less likely to experience physically incapacitating domestic violence when mills operate within the year? To answer this, I use data on monthly hospitalizations due to physical or sexual domestic violence, as hospitalization implies physical incapacitation. I show that hospitals in the catchment area of a mill are 12% and 16% less likely to have a domestic violence patient in June and July, respectively, compared to February, the month before mill operations begin. May–July marks the peak of the harvest season, when most coffee farmers in the catchment areas work in the mills. There are no statistically significant changes in domestic violence hospitalizations in January (two months before mill operations) or in the post-harvest months (August–December), when mills do not operate. While domestic violence hospitalization rates in hospitals within the catchment areas of mills remain unchanged after the harvest season relative to pre-harvest levels, they remain lower than those in hospitals outside the mills' catchment areas. These results suggest two key points. First, cases of severe domestic violence—those that lead to physical incapacitation—decline during the harvest months when mills operate. Second, part of the annual decrease in self-reported domestic violence is driven by the drop in hospitalizations during the harvest season.

As a placebo test, I examine hospitalizations for reasons other than domestic violence and find no changes, ruling out the concern that women avoid hospitals during the harvest season due to increased opportunity costs of time.

In the second part of the paper, I investigate the mechanisms behind the decline in domestic violence. A mill opening can affect violence through five channels, where the first four decrease it: First, as the wife transitions to paid work, her outside option (the utility of being single) improves, increasing her bargaining power within the marriage. Second, as household earnings increase—either through the husband's or the wife's earnings—the husband's marginal non-monetary benefit from violence (e.g., financial stress relief) declines (i.e., the income effect). Third, as the mill operates, the husband's cost of physically incapacitating his wife rises. Fourth, women's mill employment reduces couples' exposure to each other. Finally, mill wages may incentivize the husband to use violence to extract the wife's earnings, potentially increasing domestic violence.[8]

I first provide evidence that a mill opening increases the wife's bargaining power in the short run. One year after a mill opening, married women in the catchment areas score higher on a household decision-making index than the year before. They are more likely to make financial (large household purchases) and non-financial (own health, visiting family, contraception use) decisions

[7]The sexual domestic violence variable is based on women reporting that their husbands physically forced them into sex or sexual acts when they didn't want to.

[8]The Rwandan context is not ideal for testing the male backlash mechanism, given that husbands earn more than their wives (Working in the mill does not necessarily challenge the husband's breadwinner role.).

independently or jointly with their husbands, rather than having these decisions made solely by their husbands. However, this effect fades over time. Since the decline in physical domestic violence persists, this suggests that additional mechanisms beyond bargaining are at play.

Second, I test the exposure reduction mechanism using a subsample of couples whose exposure to each other plausibly remains unchanged after a mill opening. I find a decline in domestic violence even among couples with different occupations, where most wives work in agriculture and their husbands work in non-agricultural manual jobs.[9] Unlike farmer couples, these couples do not work together, meaning the mill opening does not alter the time they spend with each other during work hours. Finding a decline in violence among these couples—particularly of similar magnitude to the estimate based on the full sample—suggests that exposure reduction is not the dominant mechanism behind the results.

Using earnings data, I also find that among the couples mentioned above, a mill opening do not change the husband's earnings. This shows that domestic violence declines even in households where the increase in resources comes solely from women's paid employment rather than the husband's earnings. This finding suggests three key points. First, women's paid employment is a key driver behind the decline in domestic violence in the main sample, indicating that the increase in household resources due to the husband's earnings is not the sole mechanism. Second, since the increase in women's earnings raises household resources and coincides with a decline in violence, an overall increase in household income is also a plausible mechanism. Lastly, this mechanism check provides a case where, with the husband's earnings fixed, an increase in the wife's earnings does not lead to greater violence. This suggests that the husband is not using violence to extract the wife's earnings, meaning extractive violence is not the dominant mechanism. Does women's employment reduce domestic violence through channels beyond increased personal earnings—and thus, bargaining power—and household income?

I provide four pieces of evidence supporting that the rise in the husband's cost of physically incapacitating his wife is a key mechanism behind the decline in violence. First, I show that one month after the harvest season (August), mill-exposed households continue to have higher food consumption expenditure compared to the pre-harvest period, plausibly due to savings from selling cherries to the mill and the wife's wages from mill employment. If the seasonality of violence were solely driven by the income effect (more household resources and less financial stress), we would expect lower violence hospitalizations in both during the harvest season and August compared to the pre-harvest, when mills do not operate. However, domestic violence hospitalizations exhibit no change in August compared to the pre-harvest period. Plausibly, the wife's outside option in the catchment areas remains the same right before and after the end of the harvest season, as the husband knows his wife will have a job opportunity the following year. These suggest that the

[9]Examples include plumbers, construction workers, etc. These occupations do not change with a mill opening.

seasonal decline in domestic violence hospitalizations is driven by the husband's increased cost of physically incapacitating the wife, which peaks when mills operate, rather than by an income effect.[10]

Second, I find that in Irish potato–suitable areas, where the value of women's capacity to work in unpaid tasks does not experience a structural increase—as it does with coffee—and where women generally do not contribute to household earnings through paid work, domestic violence hospitalizations do not decline during the potato harvest months, despite higher food consumption expenditure compared to pre-harvest levels. This further confirms that when women's outside options remain constant and the value of their labor stays mostly fixed throughout the year, domestic violence hospitalizations remain unchanged.

Third, I show that self-reported emotional domestic violence (e.g., humiliation, insults)—which does not physically incapacitate women—remains unchanged after a mill opening. Combined with the earlier findings that physical domestic violence declines consistently while the decline in sexual domestic violence fades over time, this pattern suggests that the type of violence most likely to physically incapacitate women is also the only type that persistently declines, reinforcing the incapacitation channel.

Lastly, I find that gender-based violence hospitalizations for women aged 10–18, who are not the primary age group participating in the harvest or working in mills, do not exhibit a seasonal decline. These cases are primarily sexual and thus less likely to be perpetrated by the victims' fathers. Since these perpetrators lack an economic self-interest in the victim's work capacity, this finding further supports the incapacitation cost mechanism in the husband-wife relationship.

At the end of the paper, I also replicate the seasonal pattern using the universe of monthly domestic violence police records in Uganda for more than a decade which provides external validity for the incapacitation cost mechanism.

This study makes two main contributions. First, to the best of my knowledge, this is the first paper to introduce the cost of women's physical incapacitation as a channel linking women's employment to domestic violence and to provide causal evidence supporting it. A growing empirical literature on domestic violence identifies several mechanisms through which economic factors influence violence, including increases in women's outside options (Aizer 2010, Anderberg et al. 2016, Hidrobo et al. 2016, Haushofer et al. 2019, Adams-Prassl et al. 2024, Lowes 2023), reductions in financial stress (Angelucci 2008, Bhalotra et al. 2019, Heath et al. 2020, Arenas-Arroyo et al. 2021, Bhalotra et al. 2021), exposure (Chin 2011, Bhalotra et al. 2021), extractive/instrumental violence (Bloch and Rao 2002, Eswaran and Malhotra 2011, Bobonis et al. 2013, Heath 2014, Anderson and Genicot 2015, Erten and Keskin 2018, Bhalotra et al. 2019, Calvi and Keskar 2021,

[10]This does not rule out the possibility that the income effect plays a role in reducing non-severe domestic violence cases, where women do not require hospitalization and self-report in the annual household survey.

Erten and Keskin 2021, Erten and Keskin 2022), and male backlash (Angelucci 2008, Luke and Munshi 2011, Tur-Prats 2021, Alesina et al. 2020, Guarnieri and Rainer 2021). I provide evidence that while bargaining and household income channels influence domestic violence, the cost of physical incapacitation channel also plays a role—an additional mechanism that has not been systematically studied in the literature.

My results also speak to the husband's incentives behind inflicting violence. Although I did not find evidence supporting the extractive violence mechanism, the evidence supporting the incapacitation channel suggests that the husband's decision not to engage in violence is based on economic incentives and a cost-benefit analysis, rather than emotional cues or loss of control (Card and Dahl, 2011). From a wider perspective, the findings align with Becker (1968)'s economic model of crime, which posits that individuals weigh costs and benefits before committing an offense.[11]

Second, the results contribute to the literature on female labor force participation (FLFP) in developing countries. Since FLFP remains strikingly low in most developing economies, existing research focuses on its determinants and the effects of interventions aimed at increasing it. These include cultural practices, trade liberalization, residence patterns, microcredit, education, cash transfers, correcting husbands' beliefs, laws, financial inclusion, and psychosocial interventions (Alesina et al. 2013, AlAzzawi 2014, Juhn et al. 2014, Gaddis and Pieters 2017, Khanna and Pandey 2020, Angelucci et al. 2015, Banerjee et al. 2015, Keats 2018, Baird et al. 2019, Bursztyn et al. 2020, Hyland et al. 2020, Field et al. 2021, McKelway 2021). In contrast, this paper focuses on the consequences of women's employment. Heath and Jayachandran (2017) reviews the FLFP literature in developing countries and highlights that increased job availability for women influences marriage, fertility, and education decisions (Jensen 2012, Sivasankaran 2014, Heath and Mobarak 2015), children's health (Qian 2008, Atkin 2009), and bargaining power (Anderson and Eswaran 2009, Majlesi 2016). My results shed light on how interventions that increase female employment—such as access to international trade—extend beyond labor supply to improve women's physical security within the household.

As for policy implications, the findings suggest that, compared to cash transfers, jobs can be more effective in reducing domestic violence when they increase *both* women's financial independence and the economic cost of their physical incapacitation (due to cases such as couples working together or women sharing earnings with their husbands). On the other side of the coin, shocks leading to women's unemployment—such as climate change-induced weather shocks and automation—may increase violence by reducing the economic cost of physical incapacitation.

[11] My results also align with Rose (2018), who uses a kink in unemployment insurance benefits to distinguish economic incentives from time spent at work and the stress of job separation, showing that higher benefits reduce crime, including domestic violence, by altering economic incentives among unemployed ex-offenders in the state of Washington. In contrast, my paper distinguishes the husband's cost of physically incapacitating his wife from the rise in women's bargaining power and household income, behind the effect of women's employment on domestic violence.

The organization of the paper is as follows. Section 2 provides background information on the Rwandan coffee industry, rapid expansion of the mills and women's employment in the coffee value chain. Section 3 introduces data sources. Section 4 outlines the identification strategies. Section 5 present the empirical results. Section 6 uncover the mechanisms. Section 7 provides robustness checks. Section 8 concludes.

2 Institutional Context

In this section, I provide background on Rwanda's coffee industry, the rapid expansion of mills following agricultural reforms in the early 2000s, and women's employment in the coffee value chain. The discussion of the coffee industry and mill openings introduces the spatial variation (mill exposure) used in this paper. The examination of women's roles in the coffee value chain highlights the temporal variations driven by mill openings.

2.1 The Coffee Industry in Rwanda

Coffee production and processing. The first two stages of the coffee value chain, cultivation and processing, are the stages of interest throughout the paper. Both stages include labor-intensive processes which affect the quality of the end product.

A coffee tree produces coffee cherries that contain coffee beans. When the cherries ripen, they should be harvested (picked) by hand. The cherries do not ripen all at once, making harvesting labor-intensive. The harvest season lasts for approximately four to five months. It lasts from March to July in Rwanda and the peak of the harvest is from May to July.

After harvest, coffee needs to be processed. Processing is made up of three main tasks. The first task is cleaning, which is removing the outer layer (pulp) of the coffee beans. The coffee beans are dried under the sun on flat surfaces in the second task.[12] In the third task, the defective dried coffee beans are sorted out by hand based on their color and size, which is a labor-intensive process like harvesting.

There are two possible methods to process coffee cherries: dry and wet. In dry processing, also known as the traditional method due to being the oldest coffee processing method, all tasks are done by farmers at their homes without machinery.[13] The outcome of this process yields a low-quality product sold in the local market for a very low price.

Agricultural Reforms. In 2000, Paul Kagame came into power and prioritized economic growth

[12]The dried beans are called parchment coffee.

[13]Farmers clean cherries using rocks and then dry the beans on mats (Macchiavello and Morjaria, 2020).

to rebuild the country after the 1994 Genocide against the Tutsi in Rwanda. He launched the Vision 2020 program in 2000 (Boudreaux, 2011). The program outlined a list of goals that the government aimed to achieve by 2020. One of the main goals was to transform agriculture into a high-value sector. In light of this goal, the government adopted the National Coffee Strategy in 2002, which aimed to shift to high-quality, wet-processed coffee production to participate in the international specialty coffee market (Boudreaux, 2011). At the time, 90% of the Rwandan coffee was dry-processed and thus classified as low-quality (MINAGRI and MINICOM, 2008). The government also liberalized the coffee industry and removed barriers to international trade (Boudreaux, 2011).

The reforms transformed incentives in the coffee value chain. Farmers were incentivized to adopt wet-processing because it generates higher incomes (Boudreaux, 2011). Exporters competed to sell the improved coffee to foreign buyers, who were increasingly interested in Rwandan coffee due to its enhanced quality.

Expansion of the Mills. Between 2000-2006, governmental institutions collaborated with USAID, universities in Rwanda and U.S., and private sector partners under the Partnership to Enhance Agriculture in Rwanda through Linkages (PEARL) project. The project helped farmers to establish cooperatives, find loans and build mills in their communities in a few locations. After the project, farmers continued establishing cooperatives and building mills in their communities nationwide.

Figure 1 illustrates the expansion of mills in Rwanda over the last decades. The growth began after the National Coffee Strategy in 2002 and accelerated following the PEARL project, particularly between 2005 and 2010. The number of mills increased from 49 in 2005 to 197 in 2011, and surpassed 300 by 2018.

Mill Suitability. The spatial distribution of mills is influenced by the concentration of coffee trees. Figure 2 illustrates this distribution, showing that mill expansion was primarily concentrated in areas with a high number of coffee trees in 1999, prior to the adoption of the National Coffee Strategy in 2002.

This spatial pattern reflects the primary condition for mill suitability: the presence of sufficient coffee trees (Schilling and McConnell, 2004). Harvested coffee cherries must be transported to mills within a few hours to prevent rotting (Schilling and McConnell 2004, Macchiavello and Morjaria 2020). Consequently, a mill has a specific catchment area, serving farmers who produce coffee nearby.[14] In this paper, I exploit this spatial variation—within and outside of mill catchment areas—to define my main treatment variable, mill exposure.

While the presence of coffee trees is the key determinant of mill placement, broader agricultural

[14]Most farmers transport their harvest to mills on foot.

conditions also play a role in shaping Rwanda's suitability for coffee cultivation.[15] Figure A1 visualizes the FAO-GAEZ coffee suitability index, confirming that most of the country is suitable for coffee production. Arabica coffee, the dominant plant type in Rwanda, thrives at altitudes over 1,000 meters (3,000–6,000 feet) and temperatures between 15 and 24 °C (59–75 °F). Nicknamed "The Land of the Thousand Hills," Rwanda's hillsides provide the ideal conditions for cultivating coffee.

Figures 2 and A1 show that over time, the number of coffee trees increased in coffee-suitable areas that lacked trees in 1999. However, the rapid mill expansion between 2005 and 2010 was primarily concentrated in regions that already had coffee trees in 1999. This pattern arises because newly planted coffee trees require 3–5 years to produce coffee beans. In Section A4.1, I investigate the determinants of a mill opening and present supporting evidence that, conditional on agricultural conditions related to mill suitability, the timing and location of mill openings are random.

2.2 Women's Employment in the Coffee Value Chain and the Mills

The division of labor in the coffee value chain is highly gendered, a trend observed not only in Rwanda but across coffee-producing countries. According to a survey conducted by the International Trade Forum (ITC) in 15 coffee-producing member countries, including Rwanda—most of which are developing countries—women are primarily involved in the earlier, labor-intensive stages of the value chain. Table 1 presents the survey results, showing that women typically constitute 70% of the workforce in harvesting and processing, while only 10% participate in activities like in-country trading or exporting (Scholer, 2008). Drying and sorting are predominantly female-dominated tasks, whereas marketing and selling the product are male-dominated. As coffee transforms from a raw commodity into a value-added product, tasks in the coffee value chain increasingly become male-dominated (Council, 2015).

In this paper, I also use the temporal variations created by mill openings to examine how providing jobs to women affects domestic violence. Below, I outline how these variations reshape the roles and economic contributions of women and men before and after a mill opening.

In Rwanda, coffee farmers are primarily smallholder couples who own small plots for coffee production. Before a mill opening, the wife harvests coffee cherries alongside the husband, then processes them at home–a female-dominated task. The husband sells the home-processed coffee in the local market at a low price—a male-dominated task—and keeps the income as his personal earnings. Despite contributing significantly to labor-intensive stages, the wife does not earn personal income and remains financially dependent on the husband due to the gendered division of

[15] Coffee plants (*coffea*) are grown in tropical regions of Africa, South America, and Asia, collectively known as the "coffee/bean belt." The belt provides optimal conditions, such as specific ranges of temperature, rainfall, altitude, and soil characteristics.

Table 1: Women's Participation in the Coffee Value Chain in Coffee Producing Countries

Tasks in the value chain	Participation Variation (Min-Max)	Typical Participation
Fieldwork	10-90%	70%
Harvest	20-80%	70%
Sorting	20-95%	75%
In-country trading	5-50%	10%
Export	0-40%	10%
Other (certification, laboratories)	5-35%	20%

Source: International Trade Forum, 2008

Notes: "The Participation Variation (Min-Max)" column represents the range of women's participation rates across different countries, as reported in the original source. For example, in the case of fieldwork, the participation rate ranges from 10% to 90% across the 15 coffee-producing member countries surveyed by the International Trade Forum (ITC), including Rwanda—most of which are developing countries. Regarding "typical participation", the source does not specify whether it refers to the mean, median, or mode. The terminology used in this table follows the original report for consistency.

labor in the coffee value chain.

After a mill opening, the couple harvests and transports cherries to the mill, where the husband sells them for a premium price. Selling cherries to mills is significantly more profitable than selling home-processed coffee in the local market (Macchiavello and Morjaria, 2020), as mill-processed coffee is of higher quality and sold in the international coffee market.[16] Farmers outside the catchment areas continue processing coffee at home and selling it in the local market at lower prices. During the harvest months (March–July), mills demand paid labor for drying and sorting tasks that machinery cannot perform.[17] Since the tasks are female-dominated, the mill enables the wife to transition from being an unpaid family worker in her family plot to a wage worker in the mill for the same processing tasks as before and have personal earnings. Based on the setting, a mill opening affects women's employment in two ways. First, a mill opening increases the value of the wife's capacity to work on her tasks. Second, the mill allows the wife to work for pay outside of the home for her coffee processing tasks. As a result, she gains personal earnings, which she shares with the husband according to survey data.

In 2012, authors of Macchiavello and Morjaria (2020) surveyed coffee farmers who sell cherries to mills in Rwanda, where the survey results confirm that a mill opening affects women and men differently. On average, a mill buys cherries from 400 farmers, with 71% of the survey respondents being men, reinforcing that selling cherries is a male-dominated task.[18] To provide

[16] Macchiavello and Morjaria (2020) surveys Rwandan coffee farmers in 2019, asking about the relative profitability of the two processing methods. 98% of farmers reported that selling cherries to mills is more profitable than home processing.

[17] Within the harvest months on a given day, multiple farmers bring their cherries to a mill. The cherries are combined and washed with the machinery in the mills. In the following days, mills demand paid labor for the female-dominated drying and sorting tasks where women can work.

[18] Based on my data, the remaining percentage aligns with the share of women farmers who are married to non-farmer husbands or widows.

additional context, 55% of surveyed farmers are members of coffee cooperatives and 50% of the mills are owned/run by these cooperatives. Typically, if a mill is owned or run by a cooperative, it means coffee farmers in the region came together, formed a cooperative, and built the mill. Based on the data, and my time in the field, it is common for the cooperative member husband to sell coffee cherries to the mill while his wife works at the mill during the harvest season (March–July).

3 Data

This paper utilizes four main datasets that provide both spatial and temporal variations. The first dataset is a geocoded panel of mills, which enables the construction of multiple measures of mill exposure, the key treatment variable. The second dataset, the Rwandan Demographic Health Surveys (DHS), is a geocoded, repeated cross-sectional dataset at the individual and household level, providing information on self-reported labor market outcomes and domestic violence. Because both the mills and DHS data are geocoded, I can determine whether individuals are within or outside the catchment area of a mill based on GPS coordinates, creating spatial variation. I then interact this spatial variation with the number of years an individual has been in the mill's catchment area (time variation) to analyze the effects of mill openings on outcomes, using the mill's operational year and the survey year.

The third dataset, the Integrated Household Living Conditions Survey (EICV), is another repeated cross-sectional dataset at the individual and household level, providing information on earnings. While not geocoded, I define spatial variation by whether individuals reside in a district with a mill. I then interact this spatial variation with the number of years an individual has lived in a district with a mill (time variation) to analyze the effects of mill openings on earnings.

Finally, the Rwandan Hospital Management Information Systems (HMIS) dataset provides a geocoded panel of hospitals, including the universe of monthly gender-based violence hospitalizations. With both the mills and hospital data geocoded, I define whether a hospital is within the catchment area of a mill based on GPS coordinates, determining hospitals where mill-exposed individuals plausibly seek care. I then interact this spatial variation with calendar months to analyze the impact of mill operations, which occur only during the coffee harvest months, on hospitalizations, providing the second source of temporal variation in this paper.

3.1 Panel of Mills

The spatial variation in mill exposure defines the treatment group and is derived from the following sources. I create a panel of mills with GPS coordinates for each mill in Rwanda as of 2018 by harmonizing data from the Rwanda GeoPortal and Macchiavello and Morjaria (2020).

To examine the determinants of mill openings, I match mill data with several rounds (1999, 2003, 2009, 2015) of the Rwandan Coffee Census and the FAO's Global Agro-Ecological Zones (FAO-GAEZ) coffee suitability index for 1980-2010. The coffee censuses provide data on the universe of coffee trees in Rwanda at the sub-district level (approximately 55 km²), while the FAO-GAEZ index offers coffee suitability scores at a 9 km² resolution, which I aggregate to the sub-district level.

Table A1 presents summary statistics for the mill panel: 50% of mills are owned by cooperatives.[19] Mills are built in areas where there were already coffee trees and coffee farmers in 1999, prior to the National Coffee Strategy in 2002. 64% of mills are in areas with moderate to good coffee suitability according to the FAO-GAEZ index. I provide evidence that the historical number of coffee trees and the coffee suitability index are potential predictors of mill placement, and that mill placement is uncorrelated with factors affecting the evolution of female paid employment and domestic violence rates in Rwanda. Further details are in Appendix A3.1.

I construct two measures of mill exposure: residing within the catchment area of a mill (for geo-coded datasets like demographic and health survey data and hospital data) and residing in a district with a mill (for non-geo-coded earnings data). In this study, I define a mill's catchment area as the circular zone within a 4 km radius of a mill. Most mills in Rwanda are located within a radius of 3-5 km from farmers (AgriLogic, 2018). Schilling and McConnell (2004) notes that if coffee cherries are transported on foot beyond 5 km, the risk of mold or mildew increases significantly, affecting coffee quality. This suggests that 5 km represents the maximum functional distance for a mill's catchment area. According to (Macchiavello and Morjaria, 2020), on average, mill managers report catchment areas with a radius of approximately 4.5 km. Based on all this information, I construct the catchment area by creating a 4 km radius circular buffer around each mill using ArcGIS. A household cluster from the Demographic and Health Surveys (DHS) or a hospital is considered within the catchment area of a mill if its GPS coordinates fall inside the buffer. In the robustness check section, I experiment with multiple radii and find that the results remain weakly significant with a 5 km buffer and disappear at 10 km. This is consistent with the catchment area definitions found in prior reports and research.

I consider two potential control groups in my analysis using DHS data. The first control group, used to obtain the main results and study mechanisms, comprises DHS clusters outside the mills' catchment area (4 km buffer) but within the same district as the mills. I refer to this as the "within-district approach". For robustness checks, I use an alternative control group consisting of DHS clusters within the donut-shaped area between 4 and 8 km from the mills, referred to as the "donut

[19] This statistic is based on Macchiavello and Morjaria (2020), which studies mills that started operating before 2012 and includes owner type information, while the Rwanda GeoPortal data covers all mills but does not include owner type information for every mill.

approach". Figure 3 shows a visual representation of the treatment and control groups using a mill in the Nyarugenge district and DHS clusters. In both maps, the treatment group consists of DHS clusters within the the circular buffer area around the mill (triangle). In the left-hand map, representing the within-district approach, the control group consists of DHS clusters outside the circular buffer but within the same district. In the right-hand map, representing the donut approach, the control group is limited to the DHS clusters in the donut-shaped area between 4 and 8 km from the mill.

For HMIS (hospital) data, the control group consists of hospitals outside the mills' catchment areas but within the same province as the mills. This choice is driven by the limited sample size and hospital locations: there are only 48 hospitals nationwide as of 2020, and most hospitals outside a mill's catchment area are located outside the same district as the mill. Figure A2 illustrates the treated and control hospitals. Hospitals within the mill catchment areas are treated, while those outside (but within the same province) serve as controls. The hospital specification assumes that households within a mill's catchment area are more likely to visit hospitals within the catchment area, while households outside are more likely to visit hospitals outside, due to proximity.[20] Additional details on all data sources can be found in Appendix A3.

3.2 Individual and Household Level Data

I. Rwandan Demographic Health Surveys (DHS). I first examine the impact of mill openings on paid employment and self-reported domestic violence using the DHS. The DHS is a nationally representative, cross-sectional survey of individuals and households conducted in developing countries every five years. For my analysis, I use data from the 2005, 2010/11, 2014/15, and 2019/2020 cycles.[21] These surveys collect demographic and health information from women aged 15–49 and men aged 15–59. The data also include household member information, enabling me to link couples for my analysis.

The information on employment and type of earnings is collected retrospectively during individual interviews, where respondents are asked whether they were employed in the last 12 months and, if so, whether they worked for cash. This retrospective design and the timing of the interview

[20]This assumption is plausible since the cost of traveling to a distant hospital in case of a health problem is higher than visiting a closer one within the catchment area. To validate this, I create 4 km buffers around hospitals and check the percentage of control hospital buffers that do not intersect with mill buffers. A high non-intersection rate suggests that most control hospitals are far from households within mill catchment areas, reinforcing the assumption. I find that 70% of control hospital buffers do not intersect with mill buffers, suggesting households within catchment areas plausibly visit nearby hospitals, while those outside visit hospitals farther away. If a household within a mill catchment area does visit a hospital outside, my results represent a lower bound.

[21]For the 2019/2020 cycle, I restrict observations to those before March 2020 to avoid any potential effects of COVID-19.

allows DHS data to capture five harvest years: 2004, 2005, 2010, 2014, and 2019.[22] Tables A2 and A3 provide descriptive statistics for women and men (mostly women's current or former husbands) in my sample. Among women, 88% worked in the past 12 months, but only 39% earned cash. The majority of couples are self-employed farmers working in agriculture.

Domestic violence data is also collected retrospectively during women's interviews, covering experiences of physical, sexual, and emotional violence in the past 12 months, classified per World Health Organization (WHO) guidelines. Physical domestic violence behaviors include when the husband punched, kicked, dragged, tried to choke or burn the wife, or threatened or attacked her with a knife, gun, or other weapon. Sexual violence involves forced intercourse or other sexual acts the wife did not consent to. Emotional violence includes humiliation, threats, or insults directed at the wife. Among these categories, physical violence has the highest potential to physically incapacitate the wife.

Partnered women answer the questions for their current partners, while divorced or separated women answer them for their most recent partners. Since emotional violence questions are not included in every cycle, I use physical and sexual domestic violence in my main analysis and use emotional violence when studying mechanisms. According to the data, 34% of women self-reported experiencing physical or sexual domestic violence in the past 12 months.

When studying mechanisms, specifically the bargaining power channel, I also use household decision-making questions, which ask partnered women who usually makes decisions on household purchases, their own health care, visits to family or relatives, and contraception usage.

For the working-for-cash dependent variable, my sample consists of $\sim$18,500 ever-married women whose partner or most recent partner is alive ($\sim$21,000 for the working dependent variable). Approximately 13% of the women are divorced or separated. In each survey cycle, half of the households were randomly selected for the male survey, where all men aged 15–59 were interviewed. Consequently, my male sample includes $\sim$9,500 ever-married men whose partner or most recent partner is alive, with $\sim$8,000 of those men being the current husbands of the partnered women in my women's sample. I choose ever-married men whose partner or most recent partner is alive as a sample restriction to maintain consistency with the women's sample. The results are robust to using the current husband sample only. The domestic violence module was administered in households selected for the male survey. Only one woman per household was selected for the module.[23] Consequently, my final sample for domestic violence consists of $\sim$6,000 ever-married women whose partner or most recent partner is alive. For both dependent variables, polygamous unions are excluded from the sample, as they are uncommon in Rwanda, and the dynamics ana-

[22]For example, a woman interviewed in January 2015 provides retrospective responses that capture her outcomes during the coffee harvest year 2014.

[23]For instance, if a daughter is surveyed for domestic violence by her family, the mother is not surveyed for domestic violence by her partner.

lyzed in this paper focus on monogamous relationships.[24]

There is an important limitation with this data that must be discussed. To maintain respondent confidentiality, rural clusters are randomly displaced, introducing a positional error ranging from 0 to 5 kilometers. This GPS displacement can result in measurement error and bias the results (Perez-Heydrich et al., 2013). To minimize distance measurement error, I follow Perez-Heydrich et al. (2013), who recommend using a buffer distance instead of the closest distance when relying on a distance measure. Accordingly, all distances in this paper are based on a buffer distance measure: I draw a circular buffer around each mill with a specific radius and define treatment and control groups based on whether a cluster falls inside or outside the circle. Using the closest distance measure—calculating the exact distance between a mill and a cluster—would be unsuitable due to the random displacement. I revisit this issue in the robustness checks section.

II. Integrated Household Living Conditions Surveys (EICV). I examine the impact of mill openings on earnings using the EICV, a nationally representative, repeated cross-sectional individual and household-level survey. For my analysis, I use data from the 2005, 2011, 2013/14, and 2016/17 cycles, which collect demographic and socioeconomic information from households. Given that the questions I use in this dataset are not retrospective, I treat the interview years as the harvest years in the EICV.

I use the last daily labor income (cash) reported in the EICV to analyze earnings. Tables A4 and A5 provides descriptive statistics: On average, women earn less than men. For mechanisms, I also use the household food consumption module, collected monthly. Enumerators conduct ten visits, recording household spending on a variety of food items since the last visit. I aggregate these into a combined food consumption expenditure variable.

My sample for the last daily labor income variable includes approximately 6,000 ever-married women and 6,500 ever-married men whose partner or most recent partner is alive. The majority of the men are the husbands of the women. I use the same age groups as in the DHS (15-49 for women and 15-59 for men) and exclude polygamous unions. The sample excludes individuals residing in Kigali (the capital) due to its distinct, highly urbanized labor market with earnings levels differing significantly from the rest of the country, though including it in the earnings estimations yields similar results.[25]

When doing estimation using DHS or EICV, I use samples of ever-married women and ever-married men rather than married only (unless the question does not allow for it) to avoid potential

[24]Both my DHS and EICV results are robust to including polygamous couples.

[25]The confidence intervals for the insignificant estimates in some event years are larger. I should also note that although earnings levels differ greatly between Kigali and non-Kigali regions, there are no large differences in domestic violence or employment rates between these areas, making the exclusion of Kigali in the earning estimations primarily a decision based on labor market characteristics.

sample selection bias.

3.3 Administrative Hospital Level Gender-Based Violence Data

To study the incapacitation channel, I use confidential, administrative, public hospital-level, geocoded data from the Rwandan Ministry of Health (MOH) on the universe of gender-based violence hospitalizations in the country.[26] The Rwandan Health Management Information System (HMIS) data is a monthly hospital-level panel dataset on gender-based violence hospitalizations from January 2012 to December 2019.[27] These represent severe cases where the patient has been physically incapacitated.

The data captures all official gender-based violence reports in the country, documenting the number of individuals (women and men) presenting with symptoms of physical or sexual violence, disaggregated by age groups (10–18 and older than 18). However, it does not include information on patients' marital status. To construct a proxy for domestic violence, I focus on reports involving women older than 18, as 70% of women in this age group in Rwanda are married. My sample for this key dependent variable represents data from 43 public hospitals that report gender-based violence cases, out of the total 48 public hospitals in the country by 2019.

Table A6 presents descriptive statistics for the hospitals. On average, 86% of hospitals report treating at least one woman patient older than 18 with gender-based violence each month. Figure 8 illustrates the share of hospitals reporting at least one gender-based violence case involving women older than 18 in mill-exposed and non-exposed areas. Two key patterns emerge: (1) domestic violence hospitalizations are lower in mill-exposed areas compared to non-exposed areas, and (2) hospitalizations in mill-exposed areas decline when mills are operating.

Reporting to the Hospital. I fully acknowledge that hospital records rely on women reporting the cause of injury to doctors; if a woman attributes her injury to something other than gender-based violence during her hospital visit, the case may not be classified as such in the data. Yet, there are several reasons why this concern is unlikely to drive the results, and why using both hospital records and self-reported survey data is advantageous relative to relying on self-reported survey data alone.

First, although hospital records rely on reported causes, hospital admission reflects injuries se-

[26] The sample consists of 43 public hospitals in Rwanda, including district, provincial, and national referral hospitals, that recorded gender-based violence cases during the study period. These 43 hospitals provided the gender-based violence data used in this analysis, out of the 48 public hospitals in the country by 2019. The remaining five hospitals, which do not record such violence cases, are teaching hospitals or specialized hospitals (e.g., orthopedics) (MOH, 2020). More than 90% of Rwandans are covered by the government's mandatory health insurance system, with the poorest citizens benefiting from public health services free of charge (Lu et al. 2012, Koch et al. 2022).

[27] I exclude the years 2016 and 2017 due to droughts and 2015 due to inconsistencies in monthly report forms.

vere enough to require medical attention. Thus, the decision to seek care for such injuries is plausibly less discretionary than survey disclosure of domestic violence. Moreover, since the paper's contribution concerns the incapacitation mechanism, finding a decline during the harvest season using violence cases severe enough to require hospital care provides strengthening evidence that "physically incapacitating" violence is decreasing.

Second, I find a decline in domestic violence across multiple data sources. For reporting behavior alone to explain these results, rather than the mechanisms I propose, mill openings would need to simultaneously alter hospital classification of injuries in a seasonally patterned way and affect disclosure in the DHS in the years following a mill opening. In Section 5.3, I provide evidence suggesting that this is unlikely. In particular, I examine a category of non-GBV hospitalizations, injuries due to falls, for women older than 19.[28] Injuries attributed to falls could potentially include misreported GBV cases if a woman reports a GBV-related injury as a fall during her hospital visit. However, I find no seasonality in fall-related injuries, suggesting that there is no seasonal pattern in attributing GBV injuries to non-GBV causes. Finding no seasonal pattern in this non-GBV outcome, one that could potentially include misreported GBV cases, while observing a seasonal pattern in GBV cases for women older than 18 suggests that the results are unlikely to be driven by reporting.

Lastly, I show in Section 5.3 that women in mill catchment areas are not disproportionately more likely to seek help from others (including hospitals) following a mill opening relative to women outside catchment areas, suggesting that reporting behavior is not endogenous to mill openings.

4 Empirical Specification

This section outlines the identification strategies used in the paper. The first strategy (4.1 and 4.2), a difference-in-differences (DID) event study design, identifies the causal effect of a mill opening on self-reported domestic violence and labor market outcomes among individuals within or outside the catchment areas of mills. The second strategy (4.3), a different DID event study design, identifies the causal effect of the beginning of mill operations within the year, marking the onset of the coffee harvest season, on novel monthly administrative records on gender-based violence hospitalizations for women older than 18 (a proxy for domestic violence). This second strategy compares hospitals within or outside the catchment areas of mills to capture how the mill operations affect the dynamics of domestic violence hospitalizations within the year, once the mills are built.

[28] These cases are collected for women older than age 19 rather than 18, which is why I use 19 as the cutoff.

4.1 Employment and Self-Reported Domestic Violence

To investigate the impact of a mill opening on the outcomes of women and men in the sample, I estimate the specification below.

$$Y_{istk} = \beta_0 + \sum_{k=-6}^{6} Mill_{is} \times \beta_k \mathbb{1}[\tau = k] + \mathbf{X_{it}}\phi + \alpha_s + \gamma_{dt} + \varepsilon_{istk}. \quad (1)$$

As noted in Section 3.2, the Rwandan Demographic Health Survey (DHS) is a repeated cross-section collected every five years, with new respondents in each wave. Thus, the "it" subscript denotes individuals observed at a single point in time, not tracked longitudinally.

In this specification, the event of interest is the opening of a mill. The dependent variable Y_{istk} is the outcome of interest for woman i (or man), in sub-district s, in year t and at event-time k, measured in years relative to the mill's opening. $Mills_{is}$ is a binary variable coded as 1 if woman/man i resides within the catchment area of a mill and zero otherwise by the end of the sample period (February 2020).[29]

The binary variable is interacted with event-year dummies, $\mathbb{1}[\tau = k]$, to investigate the dynamic impact of a mill opening. τ denotes the event-year, where $\tau = 0$ corresponds to the mill's opening year. For $k > 0$, $\tau = k$ indicates k years after the mill's opening year; and for $k < 0$, $\tau = k$ signifies k years before the mill's opening. The omitted category is $\tau = -1$, meaning the dynamic impact of mill exposure is estimated with respect to one year prior to the mill's opening. Over the course of k years following a woman's first time being within a mill's catchment area, she may be within the catchment areas of multiple mills. k denotes the number of years relative to her being in the catchment area of that first mill. Event time years are grouped into periods (0-1, 2-3, 4-6) to ensure a balanced number of observations across years, as mills were mostly opened between 2005 and 2010, rather than evenly across every year, and both the DHS and EICV data are collected every five years, rather than annually.[30]

I control for a limited number of predetermined individual characteristics, $\mathbf{X_{it}}$, including women's religion, education, agricultural occupation, age-group fixed effects (measured in five-year age groups), year of first marriage fixed effects and interview month and interview year fixed effects.[31]

[29] As noted in the data section, the majority of women in the women's sample (ever-married women whose partner or most recent partner is alive) are married, and the majority of men in the men's sample (with the same sample restriction as the women) are the current husbands of the married women in the women's sample.

[30] For the minimum event-year prior to the mill opening, I take into account the sample size among the event-years before the mill opening, the balance across these years, and whether the minimum year before and the maximum year after the mill opening are symmetric. This rule also applies to the maximum event-year as well. The estimates reflect what I can do to the best of my ability with the data and context at hand.

[31] Agricultural occupation is included to account for sectoral heterogeneity in exposure to mill openings. Because mill openings do not affect overall labor supply or occupational composition (Table A11), including this control does not condition on a margin affected by treatment. Excluding agricultural occupation as a control leaves the magnitudes

The age-group fixed effects account for life-cycle differences, such as variations in labor market participation, marriage and fertility dynamics, and exposure to domestic violence risk at different stages of life. The results are robust to using both age-group fixed effects and year-of-birth fixed effects (measured in five-year intervals), using only year-of-birth fixed effects, and to excluding both fixed effects.[32]

The year of first marriage fixed effects control for time-varying shocks to the marriage market, such as the Rwandan Genocide against the Tutsi, which has been documented to affect husband quality, the likelihood of experiencing domestic violence and marriage stability in Rwanda (LaMattina 2017, Sanin 2024). The data records the date of the first union (marriage or cohabitation), even for divorced or separated women. As shown in the summary statistics (Table A2), 90% of the women in my sample reported having only one union in their lifetime, where the year of the first union corresponds to the year of marriage (or cohabitation). The results remain similar whether or not I include the year of first marriage fixed effects. I include interview month and interview year fixed effects to control for potential seasonal or year-specific variations in responses, given the retrospective nature of the main outcome variables. For example, a woman interviewed in January 2015 for domestic violence outcomes in the past 12 months will be responding based on her experience in the coffee harvest year 2014, whereas a woman surveyed in September 2014 may be more likely to recall the violence due to the timing of the survey relative to the harvest season. The results are robust to their exclusion.

I included sub-district fixed effects, α_s, to control for time-invariant local characteristics (observable and unobservable), such as local social norms related to women's employment, domestic violence and gender roles. I choose to include sub-district fixed effects, not catchment area fixed effects, given that a woman can be in the catchment area of multiple mills over time. Sub-district fixed effects provide a more reliable control for local factors, as they are well-defined geographical units that avoid the overlap issues in catchment areas.

District-by-year fixed effects, γ_{dt}, control for factors that vary over time and across districts, potentially influencing both mill openings and female empowerment, such as local female political participation. The proportion of women in local government varies across districts and has increased over time. Districts represent the highest tier in local government, with district councils responsible for local development programs. An increase in female representation within a district council could lead the local government to support cooperatives aiming to open mills to promote female employment. Including district-by-year fixed effects accounts for such scenarios. As for

and dynamic pattern of the estimates largely unchanged (Tables A9 and A10), with modestly lower statistical precision in the long run event-time coefficients.

[32] Year-of-birth fixed effects control for factors that vary across cohorts. For instance, older cohorts grew up in a period when Rwanda did not have pro-women laws, potentially influencing their labor market participation and attitudes toward domestic violence compared to younger cohorts.

the "year" in district-by-year fixed effects, I refer to the harvest year here (2004, 2005, 2010, 2014, and 2019) due to the retrospective nature of the main outcome variables, as outlined in the data section. Interview year fixed effects capture year-specific variations that are not accounted for by the harvest year fixed effects. I cluster standard errors at the sub-district level.

The main dependent variables for this specification are being employed, working for cash and experiencing domestic violence in the past 12 months. All main outcomes are dummy variables. For example, working for cash in the past 12 months variable takes the value one if the respondent worked for cash in the past 12 months and 0 otherwise. Domestic violence in the past 12 months variable takes the value one if a woman experienced physical or sexual domestic violence in the past 12 months and 0 otherwise. I also use dependent variables such as occupations, being divorced or separated and household decision-making for additional results and studying mechanisms.

The coefficients of interests, denoted as β_k for $k \geq 0$, represent change in outcomes for individuals exposed to a mill for the first time, compared to those unexposed within the same district, k years after a mill opening, relative to one year before the mill's operation. For $k < 0$, which corresponds represents to the years preceding the mill opening, I anticipate the β_k estimates to be small and statistically insignificant. Such findings would support the parallel trends assumption, indicating that, in the absence of treatment, the difference between the treatment and control groups remains constant over time.

4.2 Earnings

I also test the impact of a mill opening on women and men's earnings.[33] Unfortunately, the data on these variables is not geo-coded and the smallest geographical unit in the data is the district. Thus, I employ the empirical specification below:

$$Y_{idtk} = \beta_0 + \sum_{k=-7}^{5} \textit{Mill in District}_{id} \times \beta_k \mathbb{1}[\tau = k] + \mathbf{X_{it}} + \alpha_d + \gamma_t + \varepsilon_{idtk}. \tag{2}$$

As with the DHS data in Section 4.1, the Integrated Household Living Conditions Survey (EICV) is also a repeated cross-section, meaning individuals are not tracked over time across survey waves. In this specification, the event of interest is the opening of a mill in a district. The dependent variable Y_{idtk} is the last daily labor earnings for woman i (or man), in district d, in year t and at event-time k, measured in years relative to the mill's opening. $\textit{Mill in District}_{id}$ is a binary variable coded as 1 if woman/man i resides within a district where a mill has opened by the end of the sample period (October 2017) and zero otherwise.[34]

[33] I use earnings and labor income interchangeably.

[34] As mentioned in the data section, the same sample restrictions applied to the DHS data are also applied to the EICV data. Moreover, as in the DHS specification, event time years are grouped to ensure a balanced number of

The binary variable is again interacted with event-year dummies, $\mathbb{1}[\tau = k]$, to investigate the dynamic impact of a mill opening. I use the same set of pre-determined controls as in the DHS data, except for religion and the year of marriage fixed effects since the data does not have such information. α_d is the district fixed effects. γ_t is the year fixed effects.[35] As for the "year" in year fixed effects, I again refer to the harvest year, as in the DHS data. However, since the earnings question in the EICV is not retrospective for the past year, I treat the interview year as the harvest year in the EICV. I cluster standard errors at the district level.

The coefficients of interest, denoted as β_k for $k \geq 0$, represent change in earnings for individuals exposed to a mill within their district for the first time, compared to those residing in districts without a mill, measured k years after the initial mill opening and relative to one year before its operation. A district encompasses a larger geographical area than both a catchment area and a sub-district. Following the opening of the first mill, as illustrated in Figure A3, the number of mills in a district increases over time, which elevates the likelihood of individual i being exposed to a mill. Thus, I anticipate that the mill's effects become evident a few years post-opening, rather than immediately in year zero. For $k \geq 0$, I view the earnings results as complementary to DHS results rather than primary, as the dataset is not geocoded and the specification identifies effects at a broader geographical unit level. For $k < 0$, I again anticipate the β_k estimates to be small and statistically insignificant which would support the parallel trends assumption for earnings results. Based on Chen and Roth (2024), I use pseudo-likelihood poisson regression both for women's and men's earnings due to zero values, which are included in the regression.

4.3 Monthly Hospitalizations due to Gender-Based Violence

To assess the impact of mill operations on gender-based violence hospitalizations within the year and study the incapacitation channel, I estimate the following hospital-level specification:

$$Y_{hdtm} = \beta_0 + \sum_{m=1}^{12} Hospital\ in\ Mill\ Catchment_{hd} \times \beta_m \mathbb{1}[\mu = m] + \lambda_h + \sigma_m + \gamma_{pt} + \varepsilon_{hdtm}. \quad (3)$$

In this specification, the event of interest is the beginning of mill operations, marking the onset of the coffee harvest season within the year. The dependent variable Y_{hdtm} is the monthly hospitalization outcome due to gender based violence in hospital h, in district d, in year t and at event-time m, measured in months. $Hospital\ in\ Mill\ Catchment_{hd}$ is a binary variable coded as 1, if hospital h is within the catchment area of a mill during the sample period, 2012-2019 and zero otherwise. Between 2012 and 2019, only one hospital switched from being outside a mill's

observations across years.

[35] Due to the measurement of mill exposure at the district level in this specification, incorporating district-by-year fixed effects is not feasible.

catchment area to within it, following a mill opening nearby in 2017.[36] The results are robust to both correcting the treatment/control group status in the data (the approach used for the HMIS results) and dropping it from the sample. Thus, $Hospital\ in\ Mill\ Catchment_{hd}$ is a time-invariant characteristic.

One may ask how it is possible that the number of switchers is only one, given that many mills opened between 2012 and 2019, the sample period for the hospital data. This is because, although mills continued to open over the years, they were mostly concentrated in areas where mills already existed. As a result, no more than one hospital switched from being outside the catchment area of a mill to being within it. To summarize, the spatial distribution of mills, and how they expanded over time, did not create a stock of hospitals that switched from being outside the catchment area of a mill to being within it during the 2012-2019 hospitalization data period.

The mill variable is interacted with event-month dummies, $\mathbb{1}[\mu = m]$, to investigate the dynamic impact of the start of the mill's operations within the year. μ denotes the event-month, $\mu \in \{1,...,12\}$. Let β_m measure the effect of $\mu = m$ for $m \in \{1,...,12\}$, where February, $\mu = 2$, is the reference month as one month before a mill starts to operate (one month before the beginning of the coffee harvest season). So $\beta_2 = 0$ as February is the reference month and all other β_ms are estimated.

The hospital fixed effects, λ_h, control for any hospital-specific characteristic that is fixed over time like its quality and location.[37] The month fixed effects, σ_m, controls for month-specific trends. γ_{pt} is the province-by-year fixed effects. I allow year fixed effects to differ by province, one unit higher than the district.[38] This way, I am comparing the hospitals who are in the catchment area of a mill to the ones that are not, within the same province. Hospitals who are not within the catchment area of a mill within the same province constitute a more accurate control group. I cluster standard errors at the district level.

The primary dependent variable in this analysis is a binary indicator reflecting whether a hospital reported any hospitalization due to gender-based violence among women over 18. For instance, the monthly hospitalization variable is coded as 1 if a hospital recorded at least one such case in a given month, and 0 otherwise. Binary indicators are used instead of the logarithm of hospitalizations because the monthly incidence of gender-based violence cases is typically low, with a large proportion of hospitals reporting zero, one, or two cases per month.

The coefficients of interest are β_m's when $3 \leq m \leq 7$ (March-July). For $3 \leq m \leq 7$ and gen-

[36] Kabgayi District Hospital became located within a mill's catchment area following a mill opening close to the hospital in 2017.

[37] The district fixed effects are not included due to being collinear to the hospital fixed effects.

[38] Unfortunately, there is not enough observations/hospitals to have district-by-year fixed effects. Adding district-by-year fixed effects mean comparing the hospitals in the catchment area of a mill to those not within the same district. Most of the districts have only one public (district) hospital.

der based violence hospitalizations aged older than 18 category, each β_m provides the change in hospitalizations due to domestic violence within the catchment area of a mill relative to the hospitalizations that are not, within the same province, during a mill's month of operation, relative to one month before its operation. The harvest season is at its peak during $5 \leq m \leq 7$, May-July, when many coffee-producing women in the catchment areas work in the mills. Thus, I expect to see statistically significant estimates during those months.

5 Results

5.1 Employment and Self-Reported Domestic Violence

I first estimate the dynamic impact of a mill opening (the event of interest) on women's employment, type of earnings and self-reported domestic violence experiences using equation 1. I show that a mill opening increases women's paid employment and decreases their likelihood of experiencing domestic violence in the past 12 months.

Figures 4 and 5 plot the coefficients on the interaction term between the dummy for residing in the catchment area of a mill and event-year dummies. Figure 4 uses working and working for cash as the dependent variables, while Figure 5 uses different types of domestic violence (physical or sexual, physical only, sexual only) as the dependent variables. (Tables A7 and A8 report the results.)

Starting from two years after a mill opening, being in the catchment area of a mill increases women's probability of working for cash in the past 12 months by 6 percentage points (p-value= 0.02 for years 2-3 and p-value= 0.04 for years 4-6, respectively), relative to the year before the mill opening. The estimated impact represents an increase of 16% with respect to the sample mean (0.37). The probability of working in the past 12 months remains unchanged. For both dependent variables, the coefficients for the years preceding the mill opening are close to zero and statistically insignificant, providing support for the parallel trends assumption.

One year after a mill opening, being in the catchment area of a mill decreases women's probability of experiencing physical or sexual domestic violence in the past 12 months by 12 percentage points (p-value= 0.004), relative to the year before. The magnitude of the effect decreases over time and becomes statistically insignificant 4-6 years later. To explore the long-term effects, I estimate the dynamic impact of a mill opening on physical and sexual domestic violence separately. The physical domestic violence variable includes violent behaviors that could potentially lead to physical incapacitation, such as being kicked, punched, dragged, pushed, having her arm twisted, or being threatened with a knife, gun, or other weapons. The sexual domestic violence variable is based on women reporting that their husbands physically forced them into sex or sexual acts when

they didn't want to.

Starting from one year after the mill opening, being in the catchment area of a mill decreases women's probability of experiencing physical domestic violence in the past 12 months by approximately 9 percentage points consistently (p-value= 0.04 for years 0-1, p-value= 0.02 for years 2-3, and p-value= 0.05 for years 4-6, respectively). In contrast, the effect on sexual domestic violence decreases over time and becomes statistically insignificant.

This suggests that after a mill opening, domestic violence that could potentially physically incapacitate women decreases consistently over time in the years following the mill opening (0-1, 2-3, 4-6), relative to the year before. The fade-out in physical or sexual domestic violence is primarily driven by sexual domestic violence, which includes violent behaviors relatively less likely to cause physical incapacitation. In Figure A4 and Table A13, using data from the 2010, 2014, and 2019 surveys that ask about emotional violence, I also show that a mill opening does not change the incidence of emotional domestic violence in any year.

For all dependent variables related to domestic violence, the coefficients for the years preceding the mill opening are also close to zero and statistically insignificant, indicating that the pre-trends across the treatment and control groups are parallel.

As shown in Figure A5, there is no statistically significant change in men's probability of working and working for cash in the past 12 months after a mill opening. I also estimated the impact of mill exposure on occupations. It is possible that a mill opening could induce individuals to sort into the agricultural sector. Results are reported in Tables A11 and A12. Each column represents a specific occupation: managers, sales, agricultural self-employed, agricultural employees, skilled manual, and unskilled manual. Neither the women nor men exposed to a mill sort into self-employment in agriculture (i.e., becoming farmers).

5.2 Earnings

I estimate the dynamic impact of a mill opening (the event of interest) on women and men's earnings using equation 2. Using earnings data, I show that a mill opening increases both women's and men's last daily earnings.[39]

Figure A6 plots the coefficients on the interaction term between the dummy for residing in a district where a mill has opened by the end of the sample period and event-year dummies. It uses the last daily earnings of women and men as the dependent variables.

Recall that a district encompasses a larger geographical area than a mill's catchment area. Following the opening of the first mill, as illustrated in Figure A3, the number of mills in a district increases over time, which raises the likelihood of individual i being exposed to a mill. Thus, I

[39]The majority of earnings in this context are received daily based on the data. Thus, I use the last daily earnings as my dependent variable.

anticipate that the mill's effects become evident a few years after the opening, rather than immediately in year zero. As expected, the coefficients are statistically insignificant first but increasing over time. 4-5 years after a mill opening, residing in a district with a mill increases women's last daily earnings by 28 percent (p-value= 0.02) and men's last daily earnings by 64 percent (p-value= 0.01) relative to the year before the mill opening.[40] For both dependent variables, the coefficients for the years preceding the mill opening are close to zero and statistically insignificant, providing support for the parallel trends assumption.

In the data, some households in rural areas are recorded as engaging in large-scale agricultural production. Given that men earn money from selling cherries to the mill, for these households, the income from cherry sales may be reported under the household's annual income from agricultural production (it's not possible to identify which portion comes from cherry sales). Thus, I also conducted an analysis that combines men's last daily earnings with their annual income from large-scale agricultural production (divided by 365). Although this analysis is based on a specific subsample, the results provide suggestive evidence that husbands in such households experience an increase in their earnings after a mill opening, as discussed in the context (Figure A7).

5.3 Monthly Hospitalizations for Domestic Violence

Next, I estimate the dynamic impact of the beginning of a mill's operations within the year (the event of interest) on monthly hospitalizations for gender-based violence using equation 3. I show that it is less likely for a hospital in the catchment area of a mill to have a domestic violence patient during the peak of the harvest, when labor is most intensive, and mills are most active.

Figures 7 plots the coefficients on the interaction term between the dummy for the hospital being in the catchment area of a mill and event-month dummies. It uses gender-based hospitalizations for women older than 18 as the dependent variable, which is the proxy for domestic violence.

I find that it is 10 and 14 percentage points less likely for a hospital in the catchment area of a mill to have a domestic violence patient in June and July respectively (peak of the harvest), compared to one month before the mills' month of operation, February (p-value= 0.01 each). The estimated impact represents a reduction of 12% and 16% respectively, with respect to the sample mean (0.86). There is no statistically significant change in hospitalizations when mills do not operate (January, August-December).

Figure 8 visualizes the raw means for hospitalizations for domestic violence. The trend in means are also in line with the estimates. It is important to recall that the hospitalizations results

[40]The earnings measure reflects last reported daily earnings. Men may report income from sales to the mill either as last daily earnings or as income agricultural production. The survey does not include a specific question on sales to the mill, so it is not possible to determine how these transactions are recorded. I therefore also examine the effect of mill openings on income from agricultural production.

are based on the period after the most rapid expansion of the mills, post-2012. No change in hospitalizations during the post-harvest months does not mean that the hospitalizations in the catchment and non-catchment areas are the same. As seen in Figure 8, during post-harvest months, the hospitalization rates in the mill catchment areas are still lower compared to the rates outside of the mill catchment areas.

Placebo Tests. To address concerns that the results may be driven by seasonal changes in hospital visits, I conduct several placebo tests using hospitalizations that are unrelated to gender-based violence. First, I perform the same analysis for the universe of monthly hospitalizations for bone and joint disorders other than fractures.[41] Examples of such disorders include osteoarthritis, gout, rheumatoid arthritis, lupus and bursitis.[42] I report the results in Figure A8. I find no changes in hospitalizations for these disorders within the year, specifically during the harvest months relative to one month before the harvest season. Second, I perform the same analysis using the universe of monthly hospitalizations for women older than 18, capturing all possible hospitalizations in a hospital (recorded as the total number of hospitalizations, including both gender-based and non-gender-based violence cases) for the limited number of years available. If women were less likely to go to the hospital due to an increased opportunity cost during the harvest, I would expect to see a decline in hospital visits at the peak of the harvest season. I report the results in Table A19. I find no changes in hospitalizations for this age group, particularly during the harvest months relative to the month before the harvest season, suggesting that women do not systematically avoid hospitals during this period.[43]

To address concerns about women's reporting to doctors and misclassification in the hospital records, I examine injuries due to falls, which could potentially include misreported gender-based violence cases. I find no seasonality in fall-related injuries (Table A20).[44] As additional evidence on reporting behavior from the survey data, Table A21 shows that women residing in mill catchment areas are not disproportionately more likely to seek help from others (including hospitals)

[41]This category is defined by the Ministry of Health hospital classification as "Bone and joint disorders other than fractures." These cases are not recorded under the "Gender-Based Violence" section in the monthly hospital report forms. Thus, these hospitalizations are unlikely to be related to gender-based violence. I have data on these diseases for the years 2012-2015.

[42]Given the significant variation in the number of non-GBV hospitalizations per month, with the majority of hospitals reporting at least one such hospitalization, I use the logarithm of hospitalizations to better capture the distribution of these cases, rather than a binary indicator.

[43]The analysis includes the years 2016–2019 instead of 2012–2019 due to data availability. While the primary HMIS sample excludes 2016 and 2017 (drought years), I drop only 2017 here (the more severe drought year) to ensure sufficient statistical power. Although this exercise has data limitations, it provides supporting evidence that women do not systematically avoid hospitals during this period.

[44]This category is defined in the Ministry of Health hospital classification as "falls" under the "Injuries" section. These cases are not recorded under the gender-based violence section in the monthly hospital report forms. Due to data availability, this analysis uses the years 2016–2019 rather than 2012–2019.

following a mill opening, suggesting that reporting behavior is unlikely to drive the results.

Together, these findings suggest that the observed decline in gender-based violence hospitalizations is unlikely to be driven by seasonal changes in hospital visits or systematic relabeling of injuries, but instead reflects real changes in violence.

6 Mechanisms

A mill opening can influence domestic violence through five channels. First, as women transition to paid work, their outside option—the utility of being single—improves, making the threat of divorce more credible and increasing their bargaining power within marriage, which reduces the husband's violence. Second, women's mill employment reduces couples' exposure to each other, potentially lowering violence. Third, rising household earnings—whether through the husband's or wife's income—reduce the husband's marginal non-monetary benefit from violence (e.g., financial stress relief), reflecting an income effect. Fourth, mill wages may incentivize the husband to use violence to extract his wife's earnings, potentially increasing domestic violence. Finally, during mill operations, the husband's cost of physically incapacitating his wife increases, disincentivizing violence.

In the following subsections, I rule out exposure reduction and extractive violence as primary mechanisms and provide evidence for bargaining power, household income, and the cost of women's physical incapacitation. To isolate the incapacitation channel, I use monthly hospitalization data and focus on the within-year variation during the coffee harvest months, when the husband's cost of physically incapacitating his wife is high, and the other two effects are plausibly constant. I formalize these mechanisms in a conceptual framework presented in Appendix A5.

6.1 Increase in Women's Outside Options and Bargaining Power

I first investigate whether women's paid employment opportunities in the mills translate into an improvement in their bargaining power, using geo-coded data on household decision-making. I estimate the dynamic impact of a mill opening on women's participation in household decisions and provide evidence that the mill opening increases the wife's bargaining power in the short run.

I create a standardized household decision-making index using four different household decisions: decisions on large household purchases, own health, visiting her family, and being a decision-maker for contraception usage, following Kling et al. (2007). I also construct sub-indexes: a financial sub-index (decision-making on large household purchases) and a non-financial sub-index (decision-making on own health, visiting her family, and being a decision-maker for contraception usage). A higher index value indicates that the wife is usually more likely to make

decisions either alone or jointly with her husband, rather than her husband or someone else in the family usually making the decisions for her.

Figure 9 plots the coefficients on the interaction term between the dummy for residing in the catchment area of a mill and event-year dummies. The graph at the top uses the household decision-making index as the dependent variable, while the graphs at the bottom left and right use the financial and non-financial household decision-making sub-indices as dependent variables. (Table A22 reports the results.) The estimation is based on married women, as household decision-making questions are only asked to married women in the survey.

One year after a mill opening, married women in the catchment areas are 0.12 standard deviations more likely to make household decisions by themselves or jointly with their husbands (p-value = 0.009), compared to the year before. This effect is not driven solely by financial decisions. They are also 0.10 standard deviations more likely to make non-financial household decisions by themselves or jointly with their husbands one year after the mill opening (p-value = 0.025), relative to the previous year. The magnitude of the effects for both the overall index and the sub-indices decreases over time and become statistically insignificant two years after the opening.[45]

I also find that a mill opening does not affect women's likelihood of being currently divorced or separated, relative to the year before (see Figure A9 and Table A24). Recent research suggests that women exercise their right to divorce upon experiencing domestic violence in Rwanda (Sanin, 2024), making divorce a credible threat in this context. What, then, explains the fading-out of household decision-making results? Although the wife gains paid employment opportunities with the opening of the mill, the mill also increases the husband's earnings, particularly more than the wife's. Thus, the mill increases the outside options for both partners, which may offset the gains the wife experiences in the long run. Given that the decrease in physical domestic violence following a mill opening persists over the years, the results suggest that other channels, beyond the bargaining channel, may be at play.[46]

[45] Although all the indexes are constructed based on variables where the wife makes the decisions by herself or jointly with her husband, I also created an overall index in which the contraception decision is considered as one made jointly with her husband, while for the rest of the decisions, the wife makes the decisions by herself or jointly with her husband. Research finds that women in Sub-Saharan Africa may use contraception covertly to respond to their husbands' desire for more children (Ashraf et al., 2014). Results are reported in Table A23, and the estimates using this overall index are very similar to the main results.

[46] When I dropped the women who do not own land either by themselves or jointly with their husbands (I have such info only for the later data cycles), bargaining power effect in 0-1 years is more pronounced, which suggests that although it is short lived, bargaining channel is still at play.

6.2 Exposure Reduction

Another possible mechanism behind the decline in domestic violence could be the reduction in the time couples spend together.[47] If women were farming on the same plot as their husbands before the mill expansion, starting to work in the mills for pay would be a shock to both the time they spend with their husbands and their earnings. I provide evidence that reduced exposure is unlikely to drive the results by taking advantage of information on the couple's occupations and performing a subsample analysis.

I run equation 1 using a sample of women who do not share the same occupation as their current or former husbands (I excluded women where both the woman and her current or former husband work in agricultural occupations from the main sample). Self-employed married women working in agriculture (e.g., smallholder farmers), whose husbands have skilled or unskilled manual jobs (e.g., truck drivers, plumbers, refuse workers, laborers in mining, construction, and manufacturing), constitute the majority of this subsample.[48] These women have different occupations from their current or former husbands and, as a result, do not work together during working hours. Therefore, women starting work in the mills would presumably not affect the time couples spend together, as their exposure to each other remains unchanged (since there is no exposure during work hours). However, women's workplaces shift to mills, where they work for pay. If there is a statistically significant increase in the probability of women working for cash and a decline in the probability of experiencing domestic violence in the past 12 months for this subsample, it suggests that my main result—the reduction in domestic violence—is plausibly not driven by reduced exposure.

Results for women's employment and domestic violence are visualized in Figure 10 and reported in Table A25. Starting three years after a mill opening, being in the catchment area of a mill increases the probability of women working for cash in the past 12 months by 7 percentage points, relative to the year before, which corresponds to an increase of 25% relative to the sample mean (0.28). I also find that domestic violence in the past 12 months decreases by 12 percentage points, or 39% relative to the sample mean (0.31).

The results show that even among women whose exposure to their husbands is presumably unaffected by the mill expansion, there is a statistically significant decline in domestic violence. The magnitude of the effects is similar to those in the full sample, further confirming that exposure reduction is not the main driver of the decline in domestic violence.

There are two possible concerns with this subsample analysis. First, women may work long

[47] Criminologists propose the exposure reduction theory, which argues that reducing the time between a victim and her abuser can reduce the risk of violence (Dugan et al., 1999).

[48] Unfortunately, I do not have a large enough sample size to perform a balanced DID event study using wives with agricultural occupations and husbands with manual jobs only.

enough in the mills that there is a reduction in the time couples spend together, even among these couples. Although I do not have a variable that explicitly asks women about the time spent with their husbands, data shows that mill exposure does not increase women's hours worked during the coffee harvest months among women working in agriculture. The second concern is that the subsample may consist of couples with more educated husbands compared to the primary sample, making it potentially unrepresentative. This is because, in the current subsample, husbands' occupations are mostly non-agricultural jobs. However, according to the data, the mean level of husbands' education is 5.6 for this subsample and 4.4 for the full sample (with husbands working in agriculture not excluded). Based on these points, using this subsample to suggest that exposure reduction is not the dominant mechanism behind my main results remains valid.

Beyond these results, I also investigated the mean domestic violence rates for couples who work in agriculture (and presumably work together) and couples with different occupations at baseline (before a mill opening, during the rapid expansion of the mills). The rates are similar for both groups (21% and 24%, respectively). If exposure reduction had been a mechanism behind the decline in domestic violence, the data might have shown that couples with high exposure to each other have higher baseline levels of domestic violence, which is not the case.

6.3 Increase in the Household Resources

I constructed the aforementioned subsample using the earnings data, applying the same restrictions as closely as possible given the data limitations, and provide evidence that the increase in household income is a plausible mechanism behind the decline in domestic violence in the full sample.[49] Recall that in this subsample, women mostly work in the agricultural sector, while their husbands have non-agricultural jobs.

I find that among this subsample, residing in a district with a mill does not affect the husbands' earnings over time (bottom graph of Figure 10 and Table A26). I also support this result using the full sample of earnings data and a specification that interacts the dummy for couples where the wife has an agricultural occupation and the husband has a non-agricultural job, a dummy for residing in a district with a mill, and event-year dummies. For these couples, there is no statistically significant change in the husbands' earnings, with the magnitude close to zero. Results are reported in Table A27.[50]

[49]To ensure consistency with the rule I applied in the DHS (where I made sample restrictions conditional on knowing the occupation information of both partners and ex-partners), I first restrict the EICV sample to currently married men only, as I do not have the occupation information of ex-partners in the EICV data. Among these married men, I drop any man if both he and his wife have an agricultural occupation. This process results in a small reduction in the subsample size, approximately 300 divorced or separated men whose occupation is non-agricultural, for whom I do not know their former partner's occupation. The results remain similar when I include them in the subsample.

[50]For this group, no change in husbands' earnings is plausible. Although mills may demand paid labor from men with non-agricultural occupations (e.g., construction workers to build a mill, technicians to repair a mill), the expansion

Combined with the previous violence result for this subsample, the lack of an effect on husbands' earnings shows that there is a decline in violence even among couples where the increase in household resources comes solely from women's paid employment, not from an increase in husbands' earnings. Moreover, the magnitude of the decline in violence is similar to the result based on the whole sample (approximately 10 percentage points for each case). This suggests three key points. First, women's paid employment is a key driver behind the decline in domestic violence in the main sample, and the increase in household resources due to the husband's earnings is not the sole mechanism. Second, given that the increase in women's earnings raises household resources and is accompanied by a decrease in violence, an increase in household income is also a plausible mechanism behind the results. Lastly, this mechanism check provides a case where, with the husband's earnings fixed, an increase in the wife's earnings is not accompanied by a rise in violence. This suggests that the husband is not using violence to extract the wife's earnings (extractive violence is not the dominant mechanism).

6.4 Increase in the Cost of Women's Physical Incapacitation

I provide four pieces of evidence supporting that the rising cost of women's physical incapacitation is a key mechanism behind the decline in violence.

I. Seasonal Decline in Hospitalizations When the Cost of Incapacitation Peaks. For hospitals within the catchment area of a mill, gender-based violence hospitalizations for women older than 18 exhibit seasonality, with cases decreasing during the peak of the harvest season, when the cost of women's physical incapacitation is the highest for the husband. Consistent with this, pooled analysis across the peak harvest months (Figure A10) confirms a significant reduction in these hospitalizations during the peak of the harvest, reinforcing that the decline is concentrated during the period when incapacitation costs are highest.[51] What drives this seasonal decline?

Using food consumption expenditure data, I show that one month after the harvest season,

is primarily a (persistent) shock to individuals working in agriculture. It is valid to ask what women married to non-farmer men were doing before the expansion. According to DHS data, in 2004 and 2005, 59% of these women reported agricultural self-employment as their occupation before the expansion, and 61% of them were unpaid workers at that time. This suggests that they were subsistence farmers who did not use cash crops. To better understand the effect of a mill opening on them, I use crop data and find that the share of farmer women who have coffee in their family plots increases with mill expansion over the years. However, the share of farmer women who have beans in their plots does not change over time. This suggests that these women did not have coffee on their plots before, but planted coffee trees after the expansion.

[51] While May marks the beginning of the peak harvest season, the incapacitation cost mechanism is most pronounced when the majority of women in mill catchment areas are actively engaged in labor that directly contributes to household income. Given that the data come from hospitals and capture changes at the hospital level, the effect is clearest when a substantial share of women in the catchment area are affected. I thank Referee 1 for suggesting this pooled analysis, which further strengthens the empirical evidence for this mechanism.

in August, mill-exposed households continue to enjoy higher consumption compared to the pre-harvest period. This is plausibly due to savings from harvest income derived from selling cherries to the mill and the wife's wages from the mill. However, in August, domestic violence hospitalizations remain unchanged compared to the pre-harvest period, when mills do not operate. Figure 11 plots the coefficients on the interaction term between the dummy for the hospital being in the catchment area of a mill and event-month dummies. The top graph shows gender-based hospitalizations for women older than 18 as the dependent variable, while the bottom graph shows food consumption expenditure. Since this variation occurs across months rather than years, the increase in food consumption expenditure likely reflects higher food consumption rather than changes in food prices.

As I outlined in the model, as household earnings increase, the husband's marginal non-monetary benefit from violence (e.g., financial stress relief) declines. Increased food consumption in July and August suggests that this benefit is low during these months. If the seasonality of violence were driven solely by the income effect (more resources and less financial stress), I would expect lower domestic violence hospitalizations in both July and August compared to February (one month before the harvest season). However, since hospitalizations revert to their pre-harvest level in August, this indicates that the income effect alone does not explain the seasonality of violence. The key point is the timing: when the cost of incapacitation peaks, there is a noticeable dip in hospitalizations, whereas when it is not at its peak, the pattern does not show a reduction.

Plausibly, the wife's outside option in the catchment area remains the same right before and after the end of the harvest season, as the husband knows his wife will have a job opportunity the following year.[52] Since the husband sells the couple's harvest to a mill for a high premium and the wife works for pay only during the harvest season, both the value of the wife's capacity to work and the husband's cost of physically incapacitating her increase during this period. These factors suggest that the increase in the cost of women's physical incapacitation when mills operate drives the seasonal decline in violence hospitalizations, rather than solely by an income effect or changes in the wife's outside option.

II. No Seasonal Decline in Hospitalizations When the Value of Women's Work Capacity Remains Fixed. I exploit variation in the suitability of the Irish potato, which is predominantly a cash crop from which men earn money by selling the harvest during the potato harvest season in Rwanda. In areas suitable for growing Irish potatoes, although the wife can provide labor for the

[52] Although it is plausible to assume that women's outside options remain the same right before and after the harvest season, there may be cases where this assumption does not hold, such as when a woman is unable to smooth her consumption across months. Future research using high-frequency data would be valuable for investigating the dynamics of bargaining power shifts and how bargaining power responds to monthly fluctuations or the ability to smooth consumption.

husband, the value of women's physical capacity to work in unpaid tasks does not experience a structural increase—as it does within the coffee mill catchment areas—and women generally do not contribute to household earnings through paid work. This means that, in Irish potato-suitable areas, the husband's cost of physically incapacitating the wife remains mostly fixed throughout the year.

Figure 12 plots the coefficients on event-month dummies for potato-suitable districts. The top graph uses gender-based hospitalizations for women older than 18 as the dependent variable, while the bottom graph shows food consumption expenditure. According to the results, while households exhibit higher food consumption during the peak of the harvest, plausibly reflecting reduced financial stress, this does not correspond to a decrease in domestic violence hospitalizations during this period. These results support the idea that when the value of the wife's physical capacity to work remains mostly fixed throughout the year, domestic violence hospitalizations remain unchanged.

III. Consistent Decline in Self-Reported Physically Incapacitating Domestic Violence After Mill Openings. According to the self-reported domestic violence results, women consistently experience less physical domestic violence after a mill opening across all years in the analysis (0-1, 2-3, 4-6) relative to the year before the mill opening. The decline in sexual domestic violence fades over time, while a mill opening does not affect emotional domestic violence. Given the nature of violent behaviors associated with each type of violence (e.g., physical domestic violence includes acts like kicking, punching, and involving weapons), the pattern in the results suggests that the type of violence most likely to physically incapacitate the wife decreases consistently after the mill opening.

Even if the wife chooses not to work for pay in a mill in a given year, the couple still harvests the coffee cherries each harvest season. Thus, the cost of physically incapacitating the wife remains high for the husband every year after the mill opening, relative to the year before. The consistent decrease in physically incapacitating violence following the opening aligns with these facts and supports the incapacitation channel.

The consistent rate of decline in physically incapacitating violence (likely to require significant physical effort from the husband) over multiple years also supports the idea that the possibility of the husband having less energy to inflict violence due to the mill-building process is unlikely. If this were the case, we would expect a stronger effect in the 0-1 years.

IV. No Seasonal Decline in Hospitalizations When the Perpetrator Lacks Economic Self-Interest in the Victim's Work. Hospitalization data indicate that gender-based violence cases involving women aged 10–18 are primarily due to sexual violence rather than physical violence, suggesting that the perpetrator is less likely to be a household member (e.g., the father). Anecdotal

evidence also suggests that young women in this age group are unlikely to work in mills.[53] These two facts suggest that, unlike in a husband-wife relationship, perpetrators of violence against this group lack an economic self-interest in the victim's work capacity. Consistent with this, I find no seasonal decline in gender-based violence hospitalizations for women aged 10–18, further supporting the incapacitation cost mechanism that operates in the husband-wife relationship.[54]

6.5 Discussion

Mechanisms and Women's Employment. The evidence supporting the increase in the cost of women's physical incapacitation does not negate the role of increased outside options and household resources, as shown in household surveys. First, hospitalizations reflect severe domestic violence cases. Evidence for the incapacitation cost mechanism does not rule out the possibility that the income effect contributes to reducing non-severe domestic violence cases, where women do not require hospitalization and instead self-report in the annual household survey.

Second, the hospitalization data covers 2012–2019, a period during which the number of hospitals within the catchment areas of mills remained unchanged. Thus, hospitalization results capture seasonal changes within these hospitals relative to those outside the catchment areas once the mills are built, but they do not measure the effect of a mill opening, as the household surveys do. As shown in Figure 8, hospitals within the mill catchment areas already have lower domestic violence rates before the harvest season compared to those outside. This suggests that, relative to non-mill areas, women's paid employment in catchment areas increases their outside options and household resources, leading to lower domestic violence levels. By holding women's outside options constant within the year, areas with mills experience a seasonal decline in severe, physically incapacitating violence cases (hospitalizations) during the harvest season, driven by the increased cost of incapacitation.

Taken together, the evidence suggests that women's employment plays a central role in reducing domestic violence, not only by increasing household resources and outside options but also by raising the husband's economic cost of physically incapacitating his wife.[55]

[53] Harvesting coffee cherries and sorting coffee beans requires high attention to detail, which does not incentivize child labor in this context, unlike Sviatschi (2022).

[54] Given that hospitalizations for women aged older than 18 are primarily due to physical rather than sexual violence—unlike the 10–18 age group—this supports the validity of using hospitalizations for women aged older than 18 as a proxy for domestic violence between husbands and wives. This distinction aligns with the idea that perpetrators in this group have an economic self-interest in the victim's work capacity.

[55] Research suggests that employment has psychosocial benefits (Hussam et al. 2022, Hussam et al. 2023). In this context, the non-monetary value of earning personal income (e.g., confidence, dignity) and workplace social networks (e.g., coworkers in the mills) are potential benefits. Although these factors may contribute to the decline in domestic violence, my data do not allow me to quantify them. However, both channels operate through women's employment, and my findings suggest that women's employment plausibly drives the decline in domestic violence in this context. Beyond that, I provide evidence that employment reduces violence through the cost of incapacitation mechanism,

External Validity & Policy Implications. Do the results indicate that providing jobs to women decreases domestic violence in every context? Unfortunately, no. Yet, a comprehensive analysis of the mechanisms sheds light on the conditions for external validity. According to the conceptual framework, providing jobs to women reduces violence if divorce is a credible threat and the husband benefits from his wife's labor market productivity. In Rwanda, divorce is a credible threat given its laws and social norms (Sanin, 2024), and both my fieldwork and data suggest that husbands benefit from their wives' work capacity. Suppose a woman does not work with her husband on a family farm, secures a factory job, and does not share her earnings with him. Then, the husband's incentive to extract his wife's earnings may outweigh or counteract the benefits of her improved outside option. If divorce is not a credible threat, the job will not enhance her outside option in the first place. This may explain why Kotsadam and Villanger (2022) find no effect of factory jobs on physical domestic violence in Ethiopia, where divorce rates are lower and more women decide independently—rather than jointly with their husbands—how to use their earnings, compared to Rwanda.[56] Thus, while jobs may benefit women more than cash transfers, policymakers should consider whether the husband has an economic self-interest in the wife's work capacity and whether separation is a credible threat when designing policies.

Additionally, I replicate the seasonal pattern using the universe of monthly domestic violence police records in Uganda between 2011 and 2024. Exploiting variation across Arabica and Robusta growing districts where coffee mills operate only in Arabica areas, I find a similar decline in domestic violence during the harvest season, with effects reverting to pre-harvest levels, providing external validity for the incapacitation cost mechanism (Figure 13).

7 Robustness Checks

Parallel Trends. For all dependent variables in the paper, the coefficients for the years or months preceding the event of interest (mill opening or the start of mill operations within the year) are close to zero and statistically insignificant, indicating that pre-trends across the treatment and control groups are parallel. These findings support the parallel trends assumption.

Measuring the Catchment Area and Selection into the Treatment Group. I use a 4 km radius circular buffer around a mill throughout the paper to measure the catchment area and define the treatment group. As a robustness check, I experiment with multiple radii (5 km, 10 km). I show that as the buffer radius increases, incorporating more untreated DHS clusters into the treatment

extending beyond the wife's earnings and household income.

[56] These statistics are based on DHS Rwanda and Ethiopia.

group, the results fade out. Figure A11 presents these results using DHS data. These findings confirm the validity of the 4 km buffer measurement and rule out concerns that households just outside the catchment area are sorting into the treatment group.

I also estimate the results using a treatment group of DHS clusters within a 4 km radius circular buffer around the mill and a control group in the donut-shaped area between 4 and 8 km from the mills. The results closely resemble the main findings, further supporting that households just outside the treatment group do not self-select into it.

A key question, then, is why women residing just outside the catchment area are not sorting into the treatment group. There are two plausible reasons for this. First, paid jobs in the mills are in short supply, limiting employment opportunities for women outside the catchment area. Yet, could women in the catchment areas work at the mills for three days a week and hire women from outside on the fourth day using their earnings, paying them at the mill wage level? If this scenario were to occur, women outside the catchment area would secure paid employment in the mills. The answer to this question leads to the second reason why women residing just outside the catchment area are not sorting into the treatment group: the mill wage is likely lower than their reservation wage. A woman farmer outside the catchment area processes coffee at home, and her husband sells it in the domestic market. She contributes to both agricultural production and home production by growing subsistence crops (e.g., wheat for bread) for household consumption. If the mill wage does not fully compensate for her contributions at home, she chooses to continue her existing work rather than commute and seek employment at the mills.

Why do women in the catchment areas have a reservation wage lower than the mill wage, given that they choose to work in the mills according to the results? Households within the catchment area likely have a different economic structure compared to those outside. First, husbands earn more from selling cherries to the mill than from selling home-processed coffee in the local market, reducing women's need to produce food for the household. This income allows households to purchase food from the market instead of relying solely on subsistence farming. Second, their proximity to the mill means they are plausibly no longer engaging in home processing. Together, these two factors free up women's time, allowing them to work at the mills and contribute additional income to the household.

Finally, this robustness check, using multiple buffer radii, alleviates concerns about the random displacement of DHS clusters and measurement error. The results are weakly significant with a 5 km buffer and disappear at 10 km, suggesting that the estimated effect is not solely due to cluster displacement. However, measurement error remains possible. Rural DHS clusters have a positional error of up to 5 km, exceeding the 4 km buffer, meaning some clusters initially outside the catchment area may be randomly displaced into it. Despite this, I still detect an effect on women's paid employment and domestic violence with a 4 km buffer and randomly displaced

clusters. This suggests the results are a lower bound, and with precise cluster coordinates, the observed effects would likely be larger.

Selection through Marriage Survival. A potential concern is that mill openings may affect which marriages survive or dissolve, thereby changing the composition of marriages within and outside of the catchment area. If so, year of first marriage fixed effects could become endogenous. I address this concern in two ways. First, I examine whether women whose first marriage began before a mill opening are more or less likely to divorce following a mill opening. Table A31 presents these results. Because the year of marriage is observed only for the first union and divorce year is not available, I construct two subsamples. Column 1 includes women whose first marriage began before a mill opening. Column 2 further restricts the sample to women with one lifetime union, providing the closest approximation to marriages formed prior to treatment and persisting through the mill opening. In both specifications, women within and outside the catchment area are not differentially likely to divorce following a mill opening. These results suggest that differences in marriage survival do not drive the main findings. Second, Tables A32 and A33 show that the main results are robust to excluding year of marriage fixed effects. This demonstrates that the estimated effects do not rely on the inclusion of year of marriage fixed effects, even if one were concerned about potential compositional differences across marriage cohorts.

Recent Developments in the Event-Study Literature. Sun and Abraham (2020) demonstrates that in event study designs with variation in treatment timing across units, the coefficient on a given lag or lead can be contaminated from other periods. Pre-trends can arise from heterogeneity of the treatment effects as well. As a robustness check for my results, I perform the alternative estimation method proposed in Sun and Abraham (2020) that is free from such contamination.[57] All of the event study results using this method are reported in Figures A12, A13, A14, A15 and A16 and are very similar to the results in the paper.

8 Conclusion

This paper presents evidence that providing employment opportunities to women decreases domestic violence when the husband has an economic self-interest in the wife's work capacity.

I establish this relationship in two steps. First, I provide causal evidence that the rapid expansion of coffee mills in Rwanda in the 2000s increased women's paid employment, raised earnings for both spouses, and reduced domestic violence. For identification, I employ two empirical strategies using self-reported domestic violence data and novel administrative records on domes-

[57] The results are based on eventstudyinteract STATA package.

tic violence hospitalizations. Second, I investigate the mechanisms behind these results. I find that women's employment reduces domestic violence through three main channels: improving women's outside options, increasing household income, and raising the husband's cost of physically incapacitating his wife. The last channel arises because mill openings increase the value of women's work capacity—through higher returns on coffee harvests and women's wage earnings at the mill.

As for policy implications, my findings suggest that in contexts where divorce is a credible threat and husbands benefit from their wives' employment—either through the structure of employment or shared earnings—jobs may be more effective in reducing domestic violence than cash transfers. Ultimately, employment is more than just income; under the right conditions, it can reshape household power dynamics and deter domestic violence.

References

Adams-Prassl, Abi, Kristiina Huttunen, Emily Nix, and Ning Zhang, "The Dynamics of Abusive Relationships," *The Quarterly Journal of Economics*, 2024, *139*, 2135–2180.

AgriLogic, "Value Chain Analysis for the Coffee Sector in Rwanda," Technical Report 2018.

Aizer, Anna, "The gender wage gap and domestic violence," *American Economic Review*, 2010, *100*, 1847–1859.

AlAzzawi, Shireen, "Trade liberalization, industry concentration and female workers: the case of Egypt," *IZA Journal of Labor Policy*, 12 2014, *3*.

Alesina, Alberto, Benedetta Brioschi, and Eliana La Ferrara, "Violence Against Women: A Cross-cultural Analysis for Africa," *Economica*, 2020.

_ , Paola Giuliano, and Nathan Nunn, "On the origins of gender roles: Women and the plough," *Quarterly Journal of Economics*, 2013, *128*, 469–530.

Anderberg, Dan, Helmut Rainer, Jonathan Wadsworth, and Tanya Wilson, "Unemployment and Domestic Violence: Theory and Evidence," *Economic Journal*, 2016, *126*, 1947–1979.

Anderson, Siwan and Garance Genicot, "Suicide and property rights in India," *Journal of Development Economics*, 2015, *114*, 64–78.

_ and Mukesh Eswaran, "What determines female autonomy? Evidence from Bangladesh," *Journal of Development Economics*, 11 2009, *90*, 179–191.

Angelucci, Manuela, "Love on the rocks: Domestic violence and alcohol abuse in rural Mexico," *B.E. Journal of Economic Analysis and Policy*, 2008, *8*.

_ , Dean Karlan, and Jonathan Zinman, "Microcredit impacts: Evidence from a Randomized Microcredit Program Placement Experiment by Compartamos Banco," *American Economic Journal: Applied Economics*, 2015, *7*, 151–182.

Arenas-Arroyo, Esther, Daniel Fernandez-Kranz, and Natalia Nollenberger, "Intimate partner violence under forced cohabitation and economic stress: Evidence from the COVID-19 pandemic," *Journal of Public Economics*, 2 2021, *194*.

Ashraf, Nava, Erica Field, and Jean Lee, "Household bargaining and excess fertility: An experimental study in zambia," *American Economic Review*, 2014, *104*, 2210–2237.

Atkin, David, "Working for the Future: Female Factory Work and Child Health in Mexico," *Working Paper*, 2009.

Baird, Sarah, Craig McIntosh, and Berk Özler, "When the money runs out: Do cash transfers have sustained effects on human capital accumulation?," *Journal of Development Economics*, 9 2019, *140*, 169–185.

Banerjee, Abhijit, Dean Karlan, and Jonathan Zinman, "Six randomized evaluations of microcredit: Introduction and further steps," *American Economic Journal: Applied Economics*, 2015, *7*, 1–21.

Bayisenge, Jeannette, "Access to Paid Work and Women's Empowerment in Rwanda," 2010, *19*, 84–106.

Becker, Gary S., "Crime and Punishment: An Economic Approach," *Journal of Political Economy*, 1968, *76*, 169–217.

_ , "Health as Human Capital: Synthesis and Extensions," *Oxford Economic Papers*, 2007.

Bhalotra, Sonia, Diogo G.C. Britto, Paolo Pinotti, and Breno Sampaio, "Job Displacement, Unemployment Benefits and Domestic Violence," *SSRN*, 2021.

_ **, Uma Kambhampati, Samantha Rawlings, and Zahra Siddique**, "Intimate Partner Violence: The Influence of Job Opportunities for Men and Women," *The World Bank Economic Review*, 11 2019.

Bloch, Francis and Vijayendra Rao, "Terror as a bargaining instrument: A case study of dowry violence in rural India," *American Economic Review*, 2002, *92*, 1029–1043.

Bobonis, Gustavo J., Melissa González-Brenes, and Roberto Castro, "Public transfers and domestic violence: The roles of private information and spousal control," *American Economic Journal: Economic Policy*, 2013, *5*, 179–205.

Boudreaux, Karol C, *Economic liberalization in Rwanda's coffee sector: A better brew for success*

Bursztyn, Leonardo, Alessandra L. González, and David Yanagizawa-Drott, "Misperceived social norms: women working outside the home in saudi arabia," *American Economic Review*, 10 2020, *110*, 2997–3029.

Calvi, Rossella and Ajinkya Keskar, "'Til Dowry Do Us Part: Bargaining and Violence in Indian Families," 2021.

Card, David and Gordon B. Dahl, "Family violence and football: The effect of unexpected emotional cues on violent behavior," *Quarterly Journal of Economics*, 2011, *126*, 103–143.

Chen, Jiafeng and Jonathan Roth, "LOGS WITH ZEROS? SOME PROBLEMS AND SOLUTIONS," *Quarterly Journal of Economics*, 2024, *139*.

Chin, Yoo Mi, "Male backlash, bargaining, or exposure reduction?: Women's working status and physical spousal violence in India," *Journal of Population Economics*, 12 2011, *25*, 175–200.

Council, SCAA Sustainability, "The Blueprint for Gender Equality in the Coffeelands," Technical Report 2015.

Dugan, Laura, Daniel S. Nagin, and Richard Rosenfeld, "Explaining the decline in intimate partner homicide: The effects of changing domesticity, women's status, and domestic violence resources," *Homicide Studies*, 1999, *3*, 187–214.

Erten, Bilge and Pinar Keskin, "For better or for worse? Education and the prevalence of domestic violence in Turkey," *American Economic Journal: Applied Economics*, 2018, *10*, 64–105.

_ **and** _ , "Female employment and intimate partner violence: Evidence from Syrian Refugee inflows to Turkey," *Journal of Development Economics*, 2021, *150*.

_ **and** _ , "Trade-Offs? The Impact of WTO Accession on Intimate Partner Violence in Cambodia," *SSRN Electronic Journal*, 2022.

Eswaran, Mukesh and Nisha Malhotra, "Domestic violence and women's autonomy in developing countries: theory and evidence," *Canadian Journal of Economics/Revue canadienne d'économique*, 2011, *44*, 1222–1263.

Field, Erica, Rohini Pande, Natalia Rigol, Simone Schaner, and Charity Troyer Moore, "On her own account: How strengthening women's financial control impacts labor supply and gender norms," *American Economic Review*, 7 2021, *111*, 2342–2375.

Gaddis, Isis and Janneke Pieters, "The Gendered Labor Market Impacts of Trade Liberalization Evidence from Brazil," *Journal of Human Resources*, 3 2017, *52*, 457–490.

Gelles, Richard, *The Violent Home: A Study of Physical Aggression Between Husbands and Wives*, Sage Publications, 1974.

Goldin, Claudia, *Human Capital*, Springer Verlag, 2016.

Guarnieri, Eleonora and Helmut Rainer, "Colonialism and female empowerment: A two-sided legacy," *Journal of Development Economics*, 6 2021, *151*.

Haushofer, Johannes, Charlotte Ringdal, Jeremy P Shapiro, and Xiao Yu Wang, "Income Changes and Intimate Partner Violence: Evidence from Unconditional Cash Transfers in Kenya," *NBER Working paper*, 2019.

Heath, Rachel, "Women's Access to Labor Market Opportunities, Control of Household Resources, and Domestic Violence: Evidence from Bangladesh," *World Development*, 5 2014, *57*, 32–46.

_ **and Mushfiq Mobarak**, "Manufacturing growth and the lives of Bangladeshi women," *Journal of Development Economics*, 7 2015, *115*, 1–15.

_ **and Seema Jayachandran**, *The causes and consequences of increased female education and labor force participation in developing countries*, Oxford University Press, 1

_ **, Melissa Hidrobo, and Shalini Roy**, "Cash transfers, polygamy, and intimate partner violence: Experimental evidence from Mali," *Journal of Development Economics*, 3 2020, *143*.

Hidrobo, Melissa, Amber Peterman, and Lori Heise, "The effect of cash, vouchers, and food transfers on intimate partner violence: Evidence from a randomized experiment in Northern Ecuador," *American Economic Journal: Applied Economics*, 2016, *8*, 284–303.

Hussam, Reshmaan, Erin M Kelley, and Gregory Lane, "The Gendered Value of Employment," 2023.

_ **, Erin M. Kelley, Gregory Lane, and Fatima Zahra**, "The Psychosocial Value of Employment: Evidence from a Refugee Camp," *American Economic Review*, 2022, *112*.

Hyland, Marie, Simeon Djankov, and Pinelopi Koujianou Goldberg, "Gendered Laws and Women in the Workforce," *American Economic Review: Insights*, 12 2020, *2*, 475–490.

ICRW, "Evolving Men: Gender Equality Survey," Technical Report 2011.

Jensen, Robert, "Do labor market opportunities affect young women's work and family decisions? Experimental evidence from India," *Quarterly Journal of Economics*, 5 2012, *127*, 753–792.

Juhn, Chinhui, Gergely Ujhelyi, and Carolina Villegas-Sanchez, "Men, women, and machines: How trade impacts gender inequality," *Journal of Development Economics*, 2014, *106*.

Keats, Anthony, "Women's schooling, fertility, and child health outcomes: Evidence from Uganda's free primary education program," *Journal of Development Economics*, 11 2018, *135*, 142–159.

Khanna, Madhulika and Divya Pandey, "Reinforcing Gender Norms or Easing Housework Burdens? The Role of Mothers-in-Law in Determining Women's Labor Force Participation," 2020.

Kling, Jeffrey R., Jeffrey B. Liebman, and Lawrence F. Katz, "Experimental analysis of neighborhood effects," *Econometrica*, 2007, *75*, 83–119.

Koch, Rachel, Theoneste Nkurunziza, Niclas Rudolfson, Jonathan Nkurunziza, Laban Bakorimana, Holly Irasubiza, Kristin Sonderman, Robert Riviello, Bethany L. Hedt-Gauthier, Mark Shrime, and Fredrick Kateera, "Does community-based health insurance protect women from financial catastrophe after cesarean section? A prospective study from a rural hospital in Rwanda," *BMC Health Services Research*, 2022, *22*.

Kotsadam, Andreas and Espen Villanger, "Jobs and Intimate Partner Violence - Evidence from a Field Experiment in Ethiopia," *Journal of Human Resources*, 2022.

LaMattina, Giulia, "Civil conflict, domestic violence and intra-household bargaining in post-genocide Rwanda," *Journal of Development Economics*, 2017, *124*.

Lowes, Sara, "Kinship Structure and the Family: Evidence from the Matrilineal Belt," *SSRN Electronic Journal*, 2023.

Lu, Chunling, Brian Chin, Jiwon Lee Lewandowski, Paulin Basinga, Lisa R. Hirschhorn, Kenneth Hill, Megan Murray, and Agnes Binagwaho, "Towards universal health coverage: An evaluation of Rwanda Mutuelles in its first eight years," *PLoS ONE*, 2012, *7*.

Luke, Nancy and Kaivan Munshi, "Women as agents of change: Female income and mobility in India," *Journal of Development Economics*, 1 2011, *94*, 1–17.

Macchiavello, Rocco and Ameet Morjaria, "Competition and Relational Contracts in the Rwanda Coffee Chain," *The Quarterly Journal of Economics*, 2020.

Majlesi, Kaveh, "Labor market opportunities and women's decision making power within households," *Journal of Development Economics*, 3 2016, *119*, 34–47.

McKelway, Madeline, "Women's Employment in India: Intra-Household and Intra-Personal Constraints," 2021.

MINAGRI and MINICOM, "Rwanda National Coffee Strategy," Technical Report 2008.

MOH, "Rwanda Health Sector Performance Report 2017-2019," Technical Report 2020.

NISR, "Seasonal Agricultural Survey 2016," Technical Report 2016.

_ , "Seasonal Agricultural Survey 2021 Annual Report," Technical Report 2021.

Perez-Heydrich, Carolina, Joshua L Warren, Clara R Burget, and Michael E Emch, "Guidelines on the Use of DHS GPS Data. Spatial Analysis Reports No. 8," Technical Report 2013.

Qian, Nancy, "Missing women and the price of tea in China: The effect of sex-specific earnings on sex imbalance," *Quarterly Journal of Economics*, 8 2008, *123*, 1251–1285.

Rose, Evan, "The Effects of Job Loss on Crime: Evidence from Administrative Data," 2018.

Sanin, Deniz, "When Do Domestic Violence Laws Work? The Role of Social Norms," 2024.

Schilling, Michèle and Tom McConnell, "Model for Siting Coffee Washing Stations in Rwanda," 2004.

Scholer, Morten, "Women In Coffee," *International Trade Forum*, 2008, pp. 32–33.

Sivasankaran, Anitha, "Work and Women s Marriage , Fertility and Empowerment : Evidence from Textile Mill Employment," 2014.

Sun, Liyang and Sarah Abraham, "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," *Journal of Econometrics*, 2020.

Sviatschi, Maria Micaela, "Making a NARCO : Childhood Exposure to Illegal Labor Markets and Criminal Life Paths," *Econometrica*, 2022, *90*.

_ **and Iva Trako**, "Gender Violence, Enforcement, and Human Capital: Evidence from Women's Justice Centers in Peru," 2021.

Tauchen, Helen V., Ann Dryden Witte, and Sharon K. Long, "Domestic Violence: A Nonrandom Affair," *International Economic Review*, 1991, pp. 491–511.

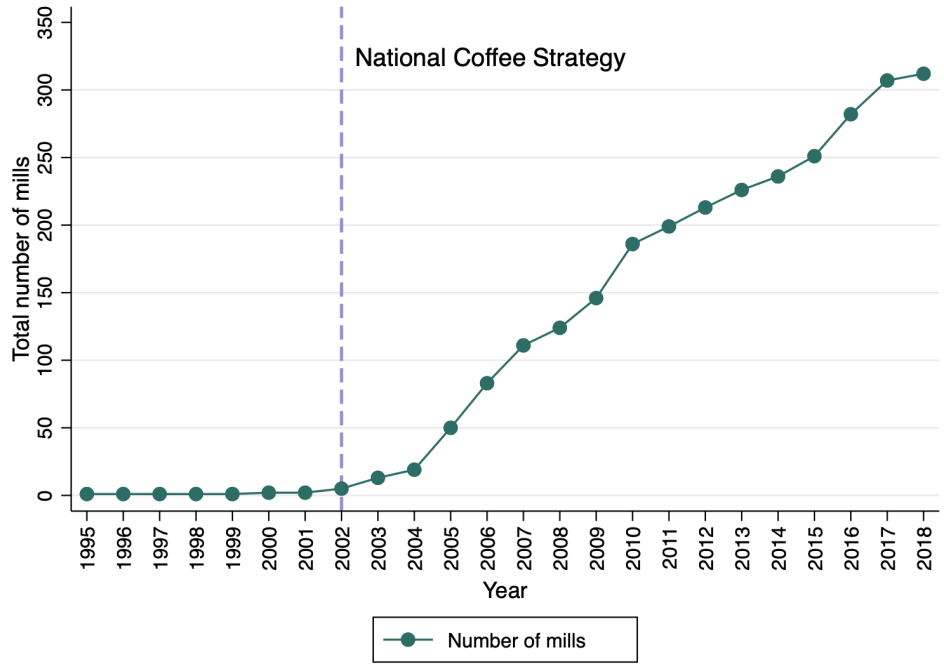
Tur-Prats, Ana, "Unemployment and intimate partner violence: A Cultural approach," *Journal of Economic Behavior and Organization*, 5 2021, *185*, 27–49.

WHO, "Global and regional estimates of violence against women: Prevalence and health effects of intimate partner violence and non-partner sexual violence," Technical Report 2013.

WorldBank, "Women, Business and the Law 2016: Getting to Equal," Technical Report, World Bank 2015.

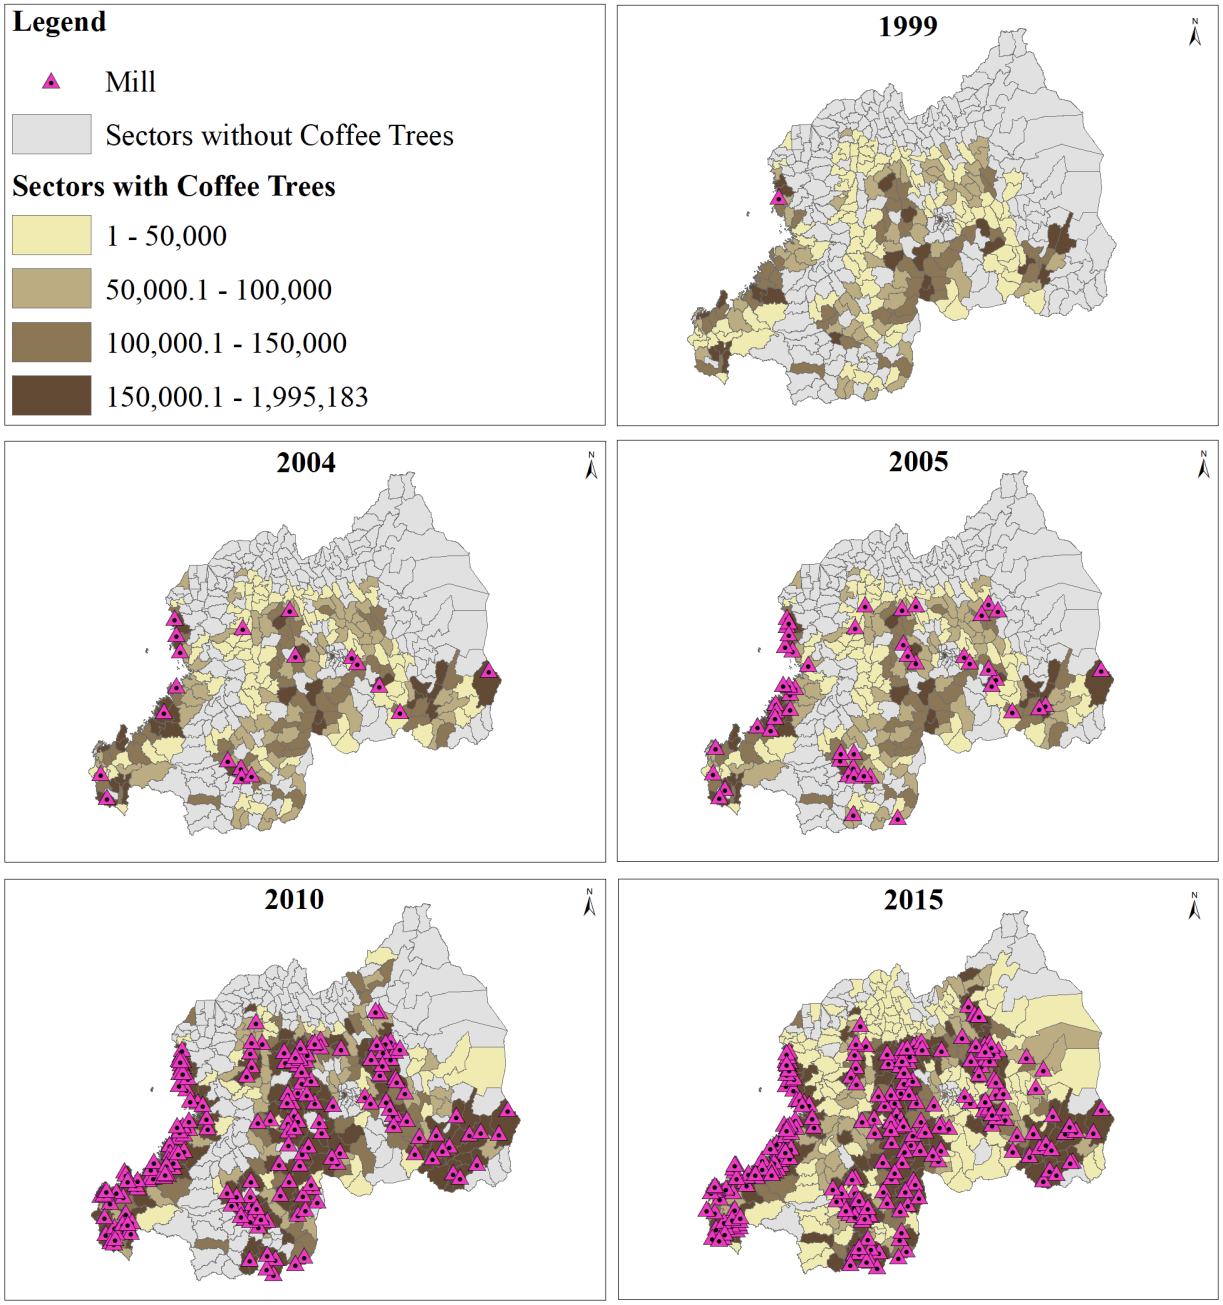
Ya-Bititi, Gisaro M., Philippe Lebailly, Deogratias Sebahire Mbonyinkebe, and Philippe Burny, *'Coffee has given us power to act': Coffee cooperatives and women's empowerment in Rwanda's rural areas: A case study of Karaba coffee cooperative*, Taylor and Francis, 8

Figure 1: Expansion of Coffee Mills in Rwanda (1995–2018)



Source: Harmonized data from the Rwanda GeoPortal, Macchiavello and Morjaria (2020), and the Rwandan Coffee Censuses (1999, 2003, 2009, and 2015).
Notes: This graph visualizes the expansion of coffee mills in Rwanda over time. The vertical dashed line marks the adoption of the National Coffee Strategy in 2002, which spurred mill expansion.

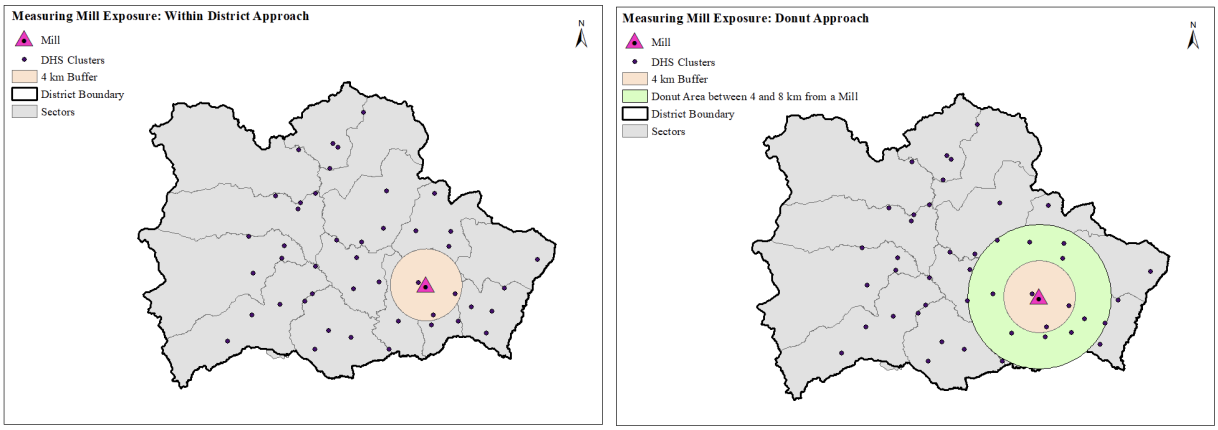
Figure 2: Mapping the Expansion of Coffee Mills in Rwanda Over Time



Source: Harmonized data from Rwanda GeoPortal, Macchiavello and Morjaria (2020), and the Rwandan Coffee Censuses (1999, 2003, 2009, 2015 rounds).

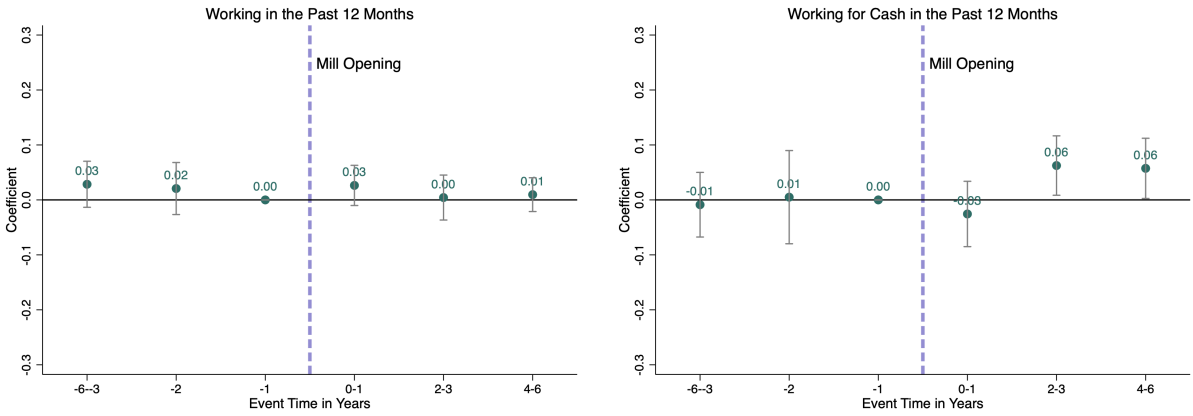
Notes: The maps display the expansion of coffee mills in Rwanda over time and across geographical units. Mills are represented as triangles, and the shaded regions indicate the number of coffee trees in each sector. Sectors without coffee trees are shown in gray. A sector is the official administrative term for a sub-district in Rwanda. While I use the term "sub-district" in the text to avoid confusion with the occupational sector, the maps retain the term "sector" to reflect the official terminology used in Rwanda. The maps are constructed using data on mills, coffee census data, and spatial sector boundaries.

Figure 3: Visualization of the Treatment and Control Groups Using Nyarugenge District: Within-District and Donut Approaches



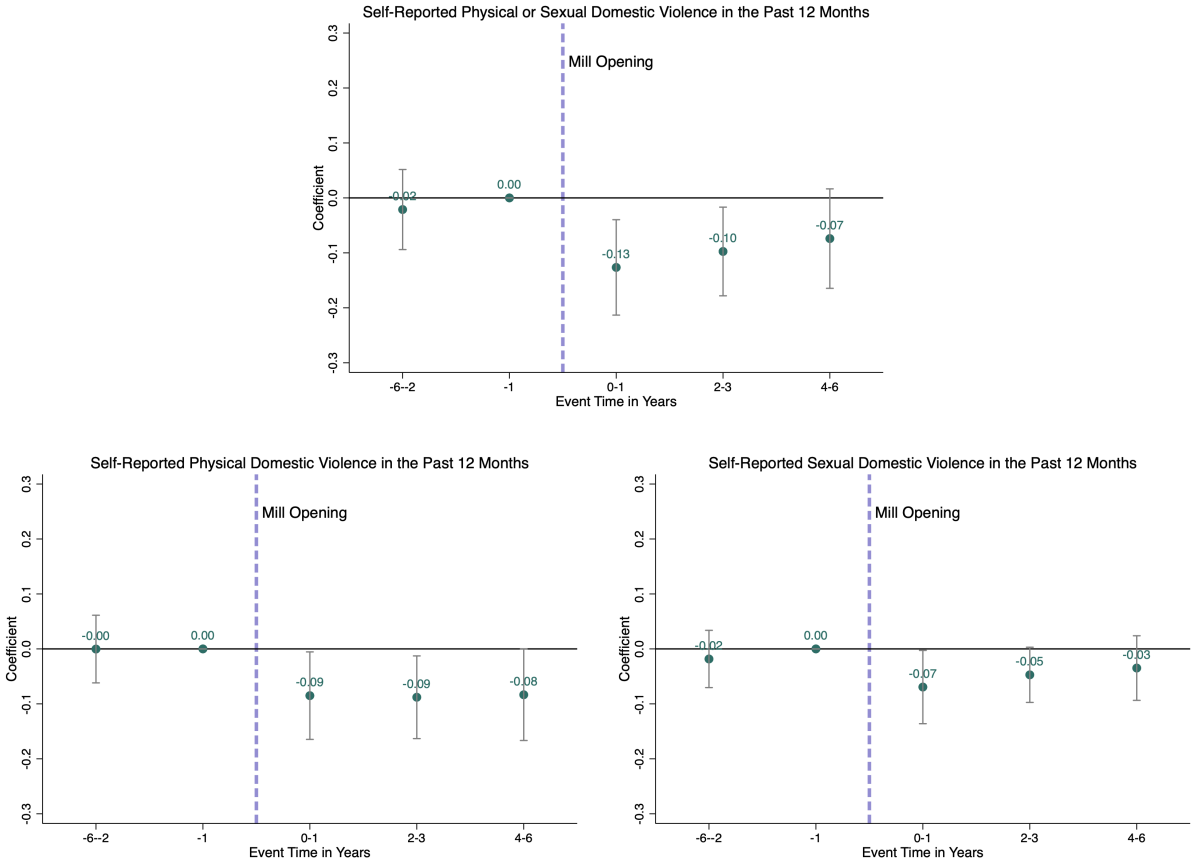
Source: Harmonized data from the Rwanda GeoPortal and Macchiavello and Morjaria (2020).
Notes: These maps visualize the treatment and control groups for a mill in the Nyarugenge district. The treatment group consists of DHS clusters within the 4 km radius circular buffer around the mill, represented by the orange circular buffer in both maps. The left map shows the "within-district approach", where the control group consists of DHS clusters outside the 4 km buffer but within the same district as the mill. The right map illustrates the "donut approach", where the control group consists of DHS clusters in the donut-shaped zone between 4 and 8 km from the mill. The within-district approach is used for the main analysis and mechanisms, while the donut approach is employed in the robustness checks section.

Figure 4: Dynamic Impact of a Mill Opening on Ever-Married Women's Working and Working for Cash in the Past 12 Months



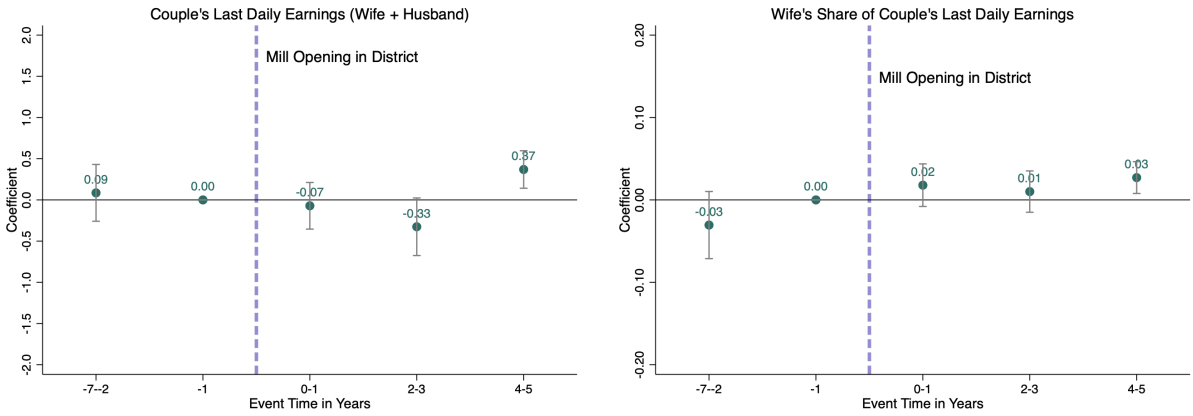
Notes: This figure plots the coefficients from the regression of women's working (left) and working for cash (right) on the interaction between residing in a mill's catchment area and event-year dummies, using equation (1). The dependent variables are drawn from the Demographic Health Surveys (DHS). The event of interest is the opening of a mill, marked by the purple dashed line. Event time is measured in years, where 0 indicates the year of the mill's opening, negative values represent years before, and positive values represent years after. Event time years are grouped to ensure a balanced number of observations across years. The sample consists of ever-married women whose partner or most recent partner is alive. The y-axis shows the change in the probability of the outcome relative to one year before the mill's opening (the omitted category). Standard errors are clustered at the sub-district level, and 95% confidence intervals are displayed.

Figure 5: Dynamic Impact of a Mill Opening on Ever-Married Women's Self-Reported Domestic Violence in the Past 12 Months: Physical or Sexual, Physical Only, Sexual Only



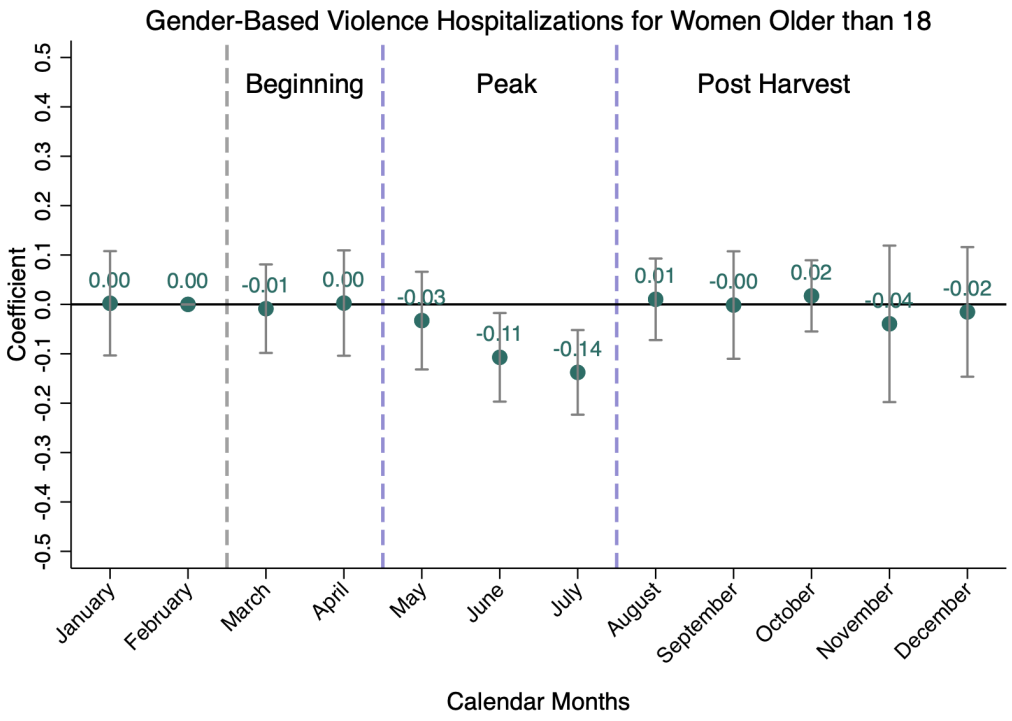
Notes: This figure plots the coefficients from the regression of self-reported domestic violence outcomes on the interaction between residing in a mill's catchment area and event-year dummies, using equation (1). The dependent variables, drawn from the Demographic Health Surveys (DHS), measure self-reported experiences of domestic violence in the past 12 months: physical or sexual violence (top), physical violence only (bottom left), and sexual violence only (bottom right). The event of interest is the opening of a mill, marked by the purple dashed line. Event time is measured in years, where 0 indicates the year of the mill's opening, negative values represent years before, and positive values represent years after. Event time years are grouped to ensure a balanced number of observations across years. The domestic violence module was administered to a randomly selected subsample of the full DHS sample. Within this subsample, I use a sample of ever-married women whose partner or most recent partner is alive. Partnered women answer the questions for their current partners, while divorced or separated women answer them for their most recent partners. The y-axis shows the change in the probability of the outcome relative to one year before the mill's opening (the omitted category). Standard errors are clustered at the sub-district level, and 95% confidence intervals are displayed.

Figure 6: Dynamic Impact of a Mill Opening on Couple's Last Daily Earnings (Wife + Husband) and the Wife's Share of Couple's Last Daily Earnings



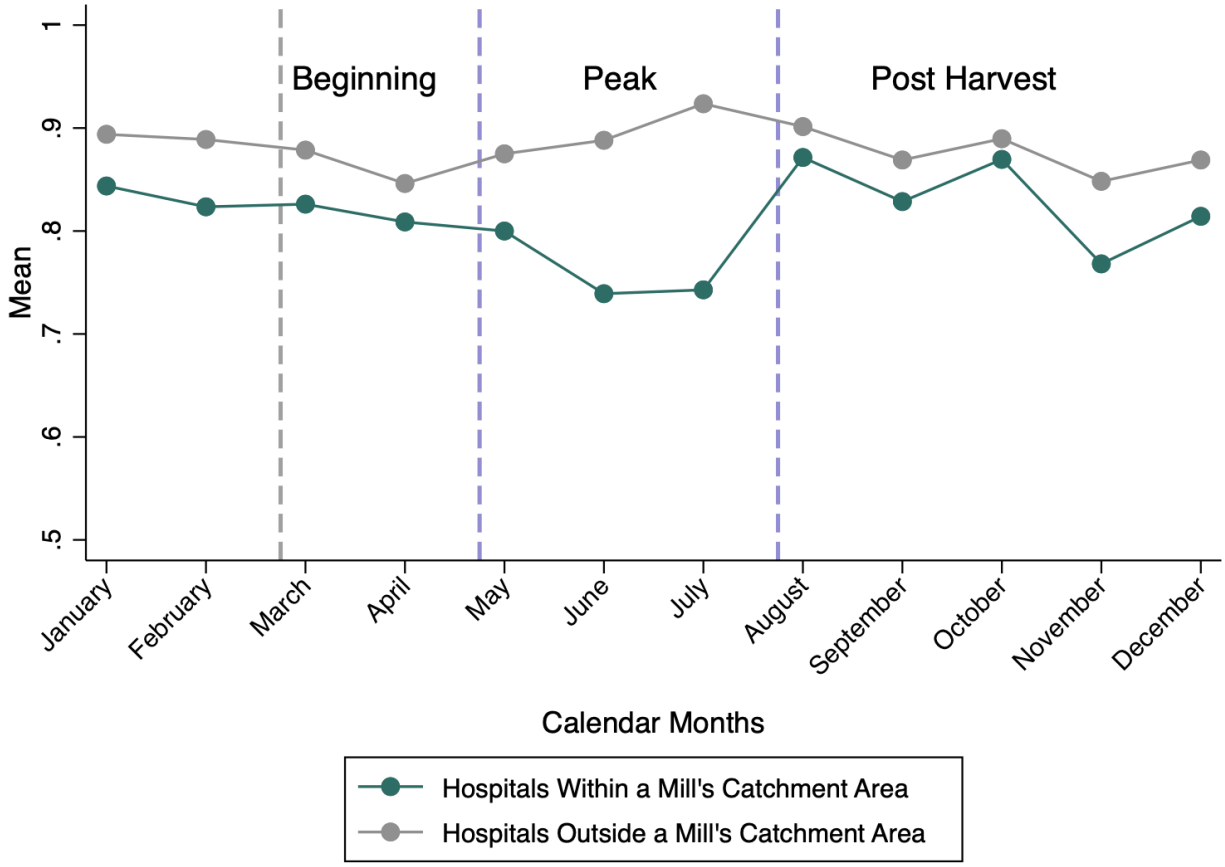
Notes: This figure plots the coefficients from the regression of couple's last daily earnings on the interaction between residing in a district with a mill and event-year dummies, using equation (2).The dependent variables, drawn from the Integrated Household Living Conditions Survey (EICV), are estimated using a Poisson pseudo-likelihood regression. The dependent variables are constructed from the Integrated Household Living Conditions Survey (EICV). The left figure examines the couple's last daily earnings, defined as the sum of the wife's and husband's last daily earnings, and is estimated using a Poisson pseudo-likelihood regression. The right figure examines the wife's share of the couple's last daily earnings, defined as the ratio of the wife's last daily earnings to total couple earnings, and is estimated using a linear regression model. The event of interest is the opening of a mill, marked by the purple dashed line. Event time is measured in years, with 0 indicating the year of the mill's opening, negative values representing years before, and positive values representing years after. Event time years are grouped to ensure a balanced number of observations across years. The sample consists of linked wife-husband pairs. The y-axis shows the change in dependent variables relative to one year before the mill's opening (the omitted category). Standard errors are clustered at the district level, and 95% confidence intervals are displayed. The sample excludes individuals residing in Kigali, the capital, due to its highly urbanized labor market, where earnings levels differ significantly from the rest of the country. Including Kigali yields similar results.

Figure 7: Dynamic Impact of the Beginning of Mill Operations on Physical or Sexual Gender-Based Violence Hospitalizations for Women Older than 18 (Proxy for Domestic Violence)



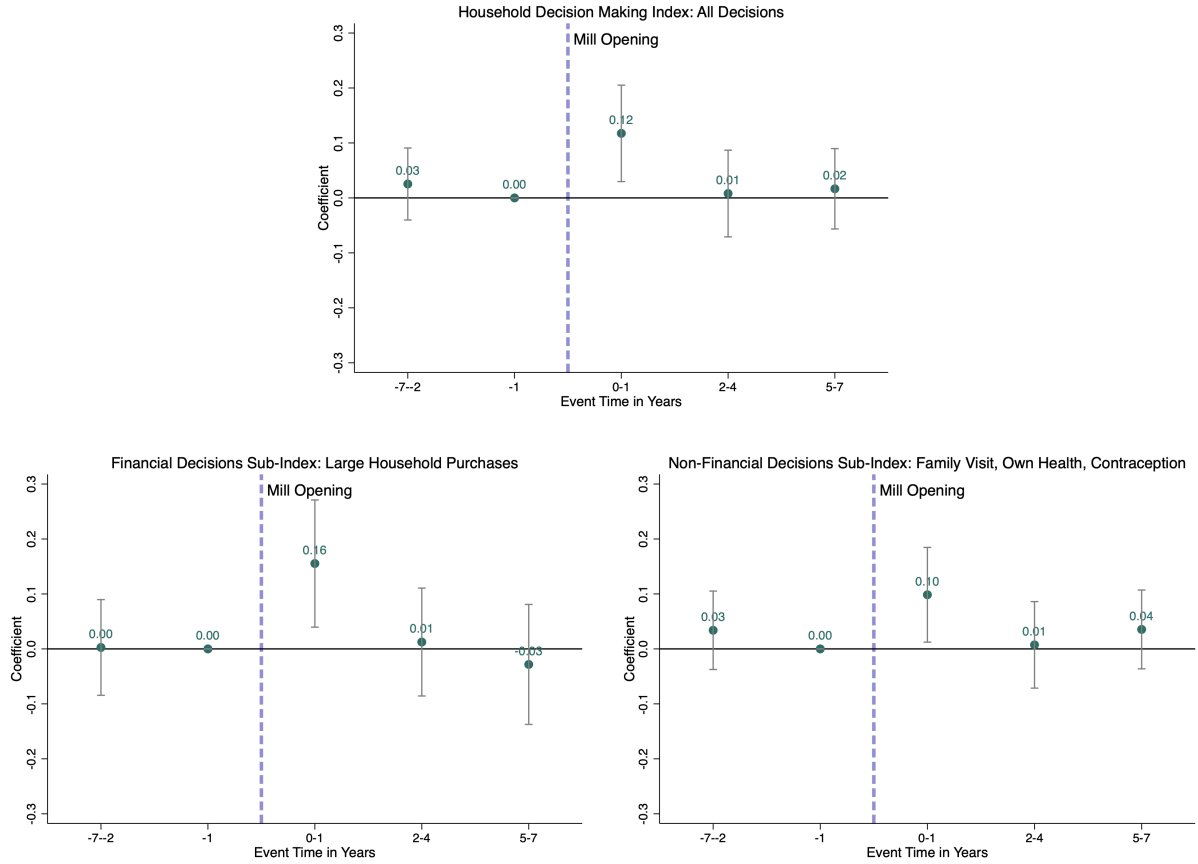
Notes: This figure plots the coefficients from the regression of hospitalizations for physical or sexual gender-based violence for women older than 18 on the interaction between a hospital being in a mill's catchment area and calendar-month dummies, using equation (3). The dependent variable, drawn from the Hospital Management Information System (HMIS) as a proxy for domestic violence, is a binary indicator reflecting whether a hospital recorded any hospitalization due to gender-based violence among women older than 18. The event of interest is the start of the mill's operations, which coincides with the beginning of the coffee harvest season, marked by the gray dashed line (March). The harvest season peaks between May and July and is followed by the post-harvest period from August to December, when mill operations have ended. Calendar months are used, with February as the reference month. The HMIS data are at the hospital level and reflect the universe of gender-based violence hospitalizations in Rwanda during the sample period. The figure specifically presents estimates for women older than 18. The y-axis shows the change in the probability of hospitalization relative to February, the month before mill operations begin. Standard errors are clustered at the district level, and 95% confidence intervals are displayed.

Figure 8: Mean Monthly Physical or Sexual Gender-Based Violence Hospitalizations for Women Older than 18 (Proxy for Domestic Violence)



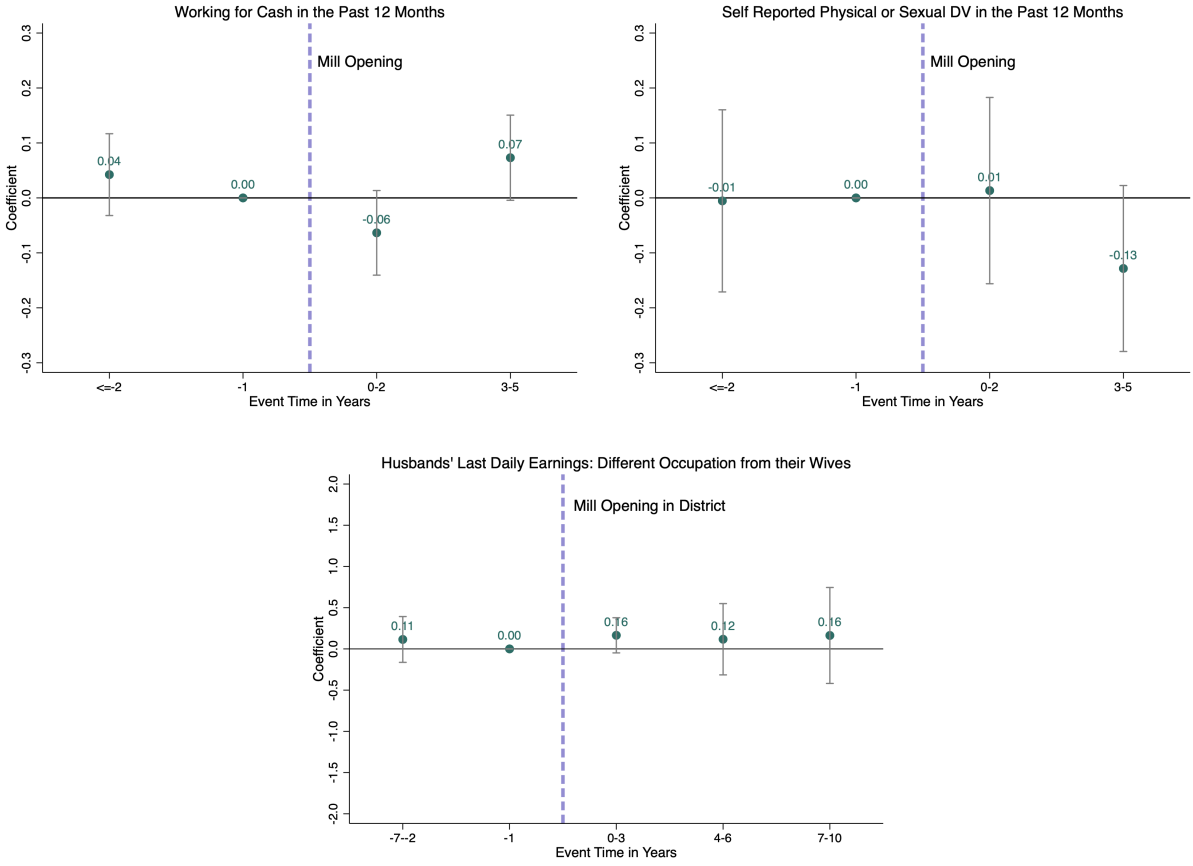
Source: Harmonized data from the Rwanda GeoPortal, Macchiavello and Morjaria (2020) and Hospital Management Information System (HMIS).
Notes: This figure shows the raw means for monthly hospitalizations due to physical or sexual gender-based violence for women older than 18. The HMIS data are at the hospital level and reflect the universe of gender-based violence hospitalizations in Rwanda during the sample period. The figure specifically presents data for women older than 18. The calendar year is divided with respect to the harvest season, during which mills operate. The gray vertical dashed line marks the start of the coffee harvest season (March). The harvest season peaks between May and July and is followed by the post-harvest period from August to December, when mill operations have ended. The green line represents hospitalizations recorded in hospitals within the catchment areas of mills, while the gray line represents hospitalizations recorded in hospitals outside the catchment areas.

Figure 9: Dynamic Impact of a Mill Opening on Married Women's Participation in Household Decision Making (Alone or Jointly with the Husband): All, Financial and Non-Financial Decisions



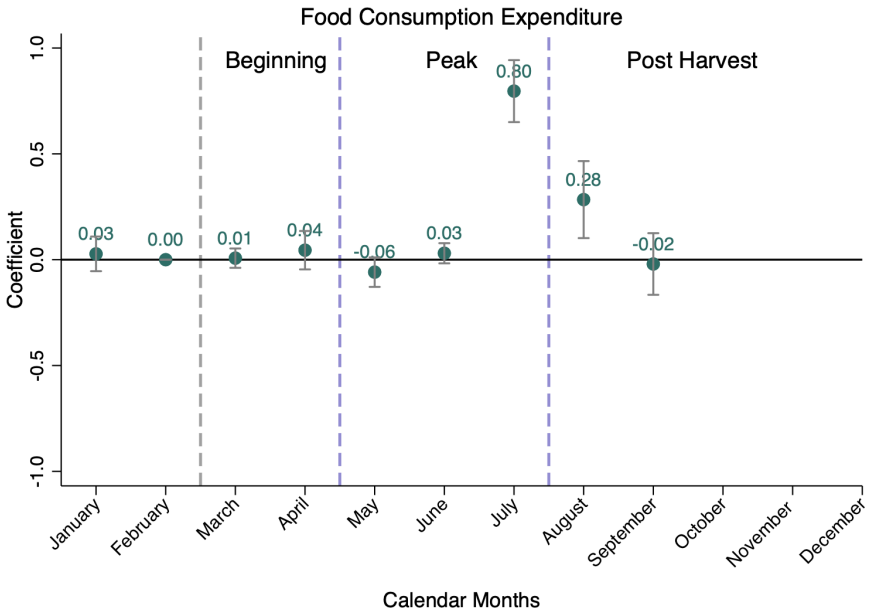
Notes: This figure plots the coefficients from the regression of household decision-making outcomes on the interaction between residing in a mill's catchment area and event-year dummies, using equation (1). The dependent variables are drawn from the Demographic Health Surveys (DHS) and are measured as standardized indices. The top figure represents an index of all household decisions where the woman reports making decisions alone or jointly with her husband. The bottom left figure focuses on financial decisions, specifically large household purchases, while the bottom right figure focuses on non-financial decisions, including family visits, own health, and contraception. The event of interest is the opening of a mill, marked by the purple dashed line. Event time is measured in years, where 0 indicates the year of the mill's opening, negative values represent years before, and positive values represent years after. Event time years are grouped to ensure a balanced number of observations across years. Since this question is collected only among couples, the sample consists of married women. The y-axis shows the change in the probability of the outcome relative to one year before the mill's opening (the omitted category). Standard errors are clustered at the sub-district level, and 95% confidence intervals are displayed.

Figure 10: Dynamic Impact of a Mill Opening on Ever-Married Women's Working for Cash and Self-Reported Domestic Violence in the Past 12 Months: Women with Different Occupations from Their Current or Most Recent Husband



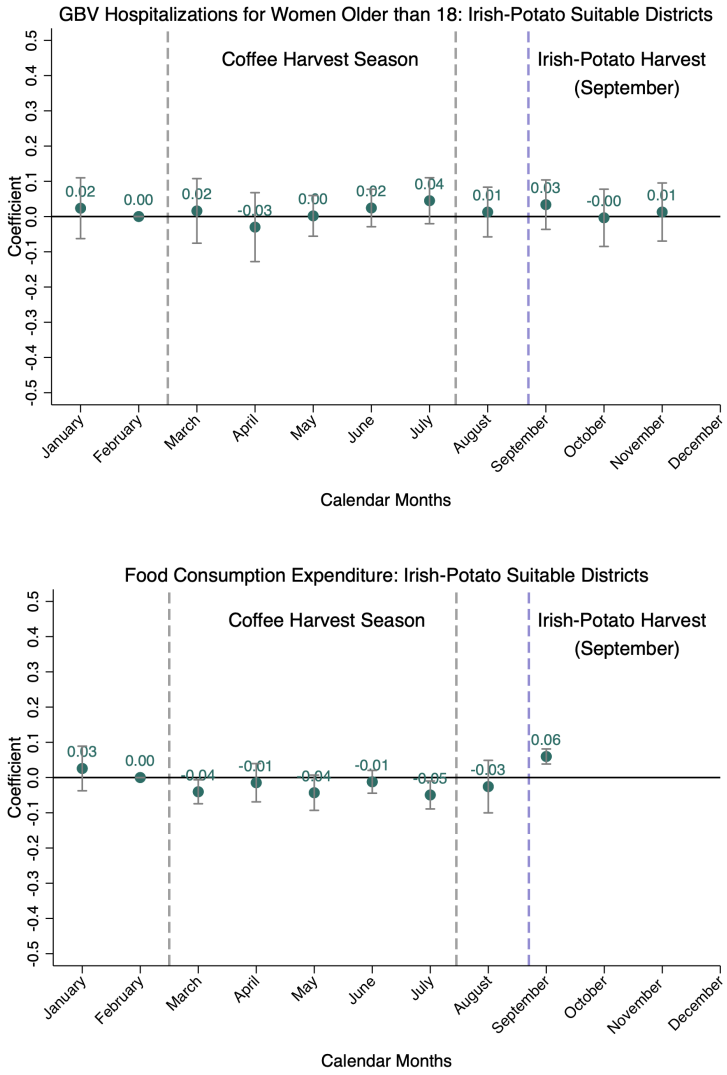
Notes: This figure plots the coefficients from the regression of women's working for cash (top-left) and self-reported physical or sexual domestic violence in the past 12 months (top-right) on the interaction between residing in a mill's catchment area and event-year dummies, using equation (1). The dependent variables are drawn from the Demographic Health Surveys (DHS). The subsample consists of ever-married women whose partner or most recent partner is alive and who do not share the same agricultural occupation as their partner (or most recent partner). I created this subsample by excluding women from the main DHS sample if both the woman and her current or most recent partner had an agricultural occupation. The bottom panel presents results for men's last daily earnings, drawn from the Integrated Household Living Conditions Survey (EICV), using equation (2). The EICV subsample is constructed to mimic the DHS subsample as closely as possible. The subsample consists of husbands of currently married women in the EICV sample who do not share the same agricultural occupation as their wives. Due to data limitations, I do not observe the most recent partner's occupation for divorced or separated men. Table A26 provides additional details on the construction of this sample and confirms that excluding the small number of divorced or separated men does not affect the results. This specification estimates earnings using a Poisson pseudo-likelihood regression. The event of interest is the opening of a mill, marked by the purple dashed line. Event time is measured in years, with 0 indicating the year of the mill's opening, negative values representing years before, and positive values representing years after. Event time years are grouped to ensure a balanced number of observations across years. The y-axis shows the change in the probability of each outcome (top panels) and earnings (bottom panel) relative to one year before the mill's opening (the omitted category). Standard errors are clustered at the sub-district level for the DHS outcomes and at the district level for the EICV outcome, and 95% confidence intervals are displayed. Kigali is not excluded in the EICV subsample, as the aim here is to create the same subsample as in the DHS analysis, and the DHS subsample does not exclude it.

Figure 11: Dynamic Impact of the Beginning of Mill Operations on Household's Monthly Food Consumption Expenditure



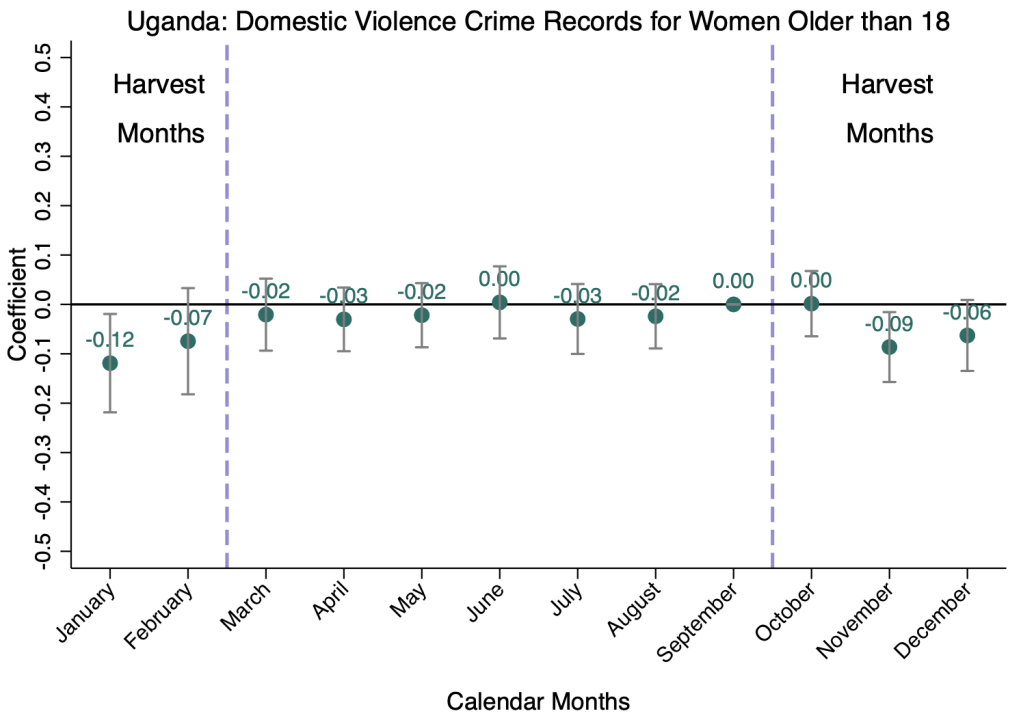
Notes: This figure plots the coefficients from the regression of household food consumption expenditure on the interaction between a district having a mill and calendar-month dummies, using an adjusted version of equation (3). The dependent variable, drawn from the Integrated Household Living Conditions Survey (EICV), is measured in logs. The event of interest is the start of mill operations, which coincides with the beginning of the coffee harvest season, marked by the gray dashed line (March). The harvest season peaks between May and July and is followed by the post-harvest period from August to December, when mill operations have ended. Calendar months are used, with February as the reference month. The sample consists of ever-married women (whose partner or most recent partner is alive) with an agricultural occupation (the majority of the full sample), as the variation in the data is tied to a harvest season. The results are consistent even without this sample restriction, with effects remaining statistically significant for July and August. The y-axis shows the change in food consumption expenditure relative to February, the month before mill operations begin. Standard errors are clustered at the district level, and 95% confidence intervals are displayed.

Figure 12: Physical or Sexual Gender-Based Violence Hospitalizations for Women Older than 18 (Proxy for Domestic Violence) and Food Consumption Expenditure in Potato-Suitable Districts



Notes: This figure plots the coefficients from the regression of gender-based violence hospitalizations for women older than 18 (top) and household food consumption expenditure (bottom) on calendar-month dummies in Irish potato-suitable districts. The hospital sample consists of hospitals located in Irish potato-suitable districts. The consumption sample consists of ever-married women (whose partner or most recent partner is alive) with an agricultural occupation residing in these districts. As in the previous consumption analysis, I restricted the sample to women with an agricultural occupation, as the variation in the data is tied to a harvest season. The dependent variables are defined as in Figures 7 and 11, restricted to this subsample. The vertical dashed lines indicate key agricultural periods: the gray dashed line marks the coffee harvest season, while the purple dashed line marks the Irish potato harvest in September. The Rwanda Seasonal Agricultural Surveys define Season C (July–September) as one of Rwanda's agricultural seasons, with Irish potatoes as one of its main crops and September as the harvest month (NISR 2016, NISR 2021). Calendar months are used, with February as the reference month. Standard errors are clustered at the district level, and 95% confidence intervals are displayed.

Figure 13: Dynamic Impact of the Beginning of the Harvest Season on Domestic Violence Crime Records for Adult Women in Uganda



Notes: This figure plots the coefficients from the regression of police reports for domestic violence for adult women on the interaction between a police station being in a Arabica coffee suitable district (where mills operate) and calendar-month dummies, using equation (3). The dependent variable, drawn from the Ugandan Police Force Crime Records, as police reports for domestic violence, is a binary indicator reflecting whether a police station recorded any police report due to domestic violence among adult women. The event of interest is the start of the mill's operations, which coincides with the beginning of the coffee harvest season, marked by the gray dashed line (October). The harvest season peaks between November and January and is followed by the post-harvest period from March to September, when mill operations have ended. Calendar months are used, with September as the reference month. The data are at the police station level and reflect the universe of domestic violence crime records in Uganda during 2011-2024. The figure specifically presents estimates for adult women. The y-axis shows the change in the probability of reporting to police relative to September, the month before mill operations begin. Standard errors are clustered at the district level, and 95% confidence intervals are displayed.

Online Appendix

for "Women's Employment, Husbands' Economic Self-Interest and Domestic Violence"

by Deniz Sanin

A1 Additional Tables **2**

A1.1 Additional Tables for Section 3: Data 2

A1.2 Additional Tables for Section 5: Results 8

A1.3 Additional Tables for Section 6: Mechanisms 24

A1.4 Additional Tables for Section 7: Robustness Checks 33

A2 Additional Figures **36**

A2.1 Additional Figures for Section 2: Institutional Context 36

A2.2 Additional Figures for Section 3: Data 37

A2.3 Additional Figures for Section 4: Empirical Specification 38

A2.4 Additional Figures for Section 5: Results 39

A2.5 Additional Figures for Section 7: Mechanisms 44

A2.6 Additional Figures for Section 8: Robustness Checks 46

A3 Data Sources **53**

A3.1 Panel of Mills 53

A3.2 Individuals and Household Level Data 53

A3.3 Administrative Hospital Level Gender-Based Violence Data 54

A4 Supplementary Empirical Analysis **55**

A4.1 What predicts mill placement? 55

A5 Theory Appendix **58**

A5.1 Theoretical Framework 58

A5.2 Setup 58

A5.3 Proofs 62

A5.3.1 Increase in I_w 62

A5.3.2 Increase in I_h 64

A1 Additional Tables

A1.1 Additional Tables for Section 3: Data

Table A1: Summary Statistics for Mills: Panel of Mill Data

	Mean	Std. Dev.
Panel A: Mill Level Characteristics		
Owner: Cooperative	0.50	0.50
Owner: NGO	0.25	0.43
Opened after National Coffee Strategy in 2002	0.99	0.08
Opened between 2005-2010	0.54	0.50
Opened between 2010-2018	0.40	0.49
Panel B: Sub-District Level Characteristics		
Number of coffee farmers in 1999	696.54	437.48
Log Coffee Trees in 1999	11.50	0.56
Log Coffee Trees in 2003	11.59	0.60
Log Coffee Trees in 2009	12.90	0.77
Log Coffee Trees in 2015	13.12	0.94
FAO-GAEZ Coffee Suitability Index for Coffee: Marginal	0.36	0.48
FAO-GAEZ Coffee Suitability Index for Coffee: Good, Medium or Moderate	0.64	0.48
Panel C: Spatial Characteristics		
Sub-District Area in km^2	54.57	26.32
District Area in km^2	877.11	304.92

Notes: This table presents summary statistics for mills, including mill-level characteristics, sub-district level characteristics, and spatial characteristics. Mill-level data are sourced from the Rwanda GeoPortal and Macchiavello and Morjaria (2020), while coffee tree and farmer data are from the national coffee censuses. The FAO-GAEZ Suitability Index for coffee categorizes land suitability into the following: Very high, High, Good, Medium, Moderate, Marginal, Very marginal, and Not suitable. The continuous index is calculated at a 9 km^2 resolution and aggregated at the sub-district level. Spatial characteristics are derived using the Rwanda GeoPortal via ArcGIS software.

Table A2: Summary Statistics for Women: Demographic Health Survey (DHS)

	All		Outside a Mill's CA		Within a Mill's CA	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Panel A: Main dependent variables						
Worked in the past 12 months	0.88	0.33	0.88	0.33	0.89	0.31
Worked for cash in the past 12 months	0.37	0.48	0.35	0.48	0.47	0.50
Experienced domestic violence in the past 12 months	0.31	0.46	0.31	0.46	0.35	0.48
Panel B: Characteristics (Used as Controls)						
Christian	0.96	0.19	0.96	0.19	0.97	0.16
Age (Controlled via age group fixed effects)	32.79	7.76	32.65	7.70	33.08	7.87
Occupation: Agricultural self-employed	0.74	0.44	0.70	0.46	0.84	0.36
Education: No education	0.20	0.40	0.20	0.40	0.18	0.39
Education: Primary not completed	0.49	0.50	0.48	0.50	0.54	0.50
Education: Primary completed	0.18	0.38	0.18	0.38	0.19	0.39
Education: More than primary	0.13	0.34	0.15	0.35	0.09	0.29
Year of first marriage (Controlled via fixed effects)	1999.51	9.59	2000.29	9.49	1998.37	9.19
Panel C: Other Characteristics (Not Used as Controls)						
Marital status: Married	0.59	0.49	0.58	0.49	0.64	0.48
Marital status: Living together	0.28	0.45	0.30	0.46	0.24	0.43
Marital status: Divorced	0.05	0.23	0.05	0.23	0.06	0.23
Marital status: Not living together	0.07	0.26	0.07	0.26	0.06	0.24
Age at first marriage	20.78	3.73	20.76	3.76	20.99	3.71
Number of unions: One	0.90	0.30	0.90	0.30	0.90	0.30
Total children ever born	3.64	2.46	3.59	2.44	3.70	2.47
Observations	21087		15925		3244	

Notes: This table presents summary statistics for women in my sample from the Demographic Health Survey (DHS). Panel A reports the main dependent variables, while Panel B includes control variables (women's characteristics) used in the specification associated with the dataset. Age is controlled via age group fixed effects, measured in five-year age groups. Year of first marriage fixed effects (measured in years and collected for ever-married women) are included to account for shocks to the marriage market. I also include (women's) interview month and interview year fixed effects in the specification, where the results are robust to their exclusion. Panel C provides additional characteristics related to marital status and family dynamics (not used as controls). The sample is split by women living within or outside the catchment area of a mill (Mill's CA), which is defined as a 4 km radius buffer around a mill. The domestic violence module was administered to a randomly selected subsample of the full DHS sample, and thus, my sample for the domestic violence outcomes consists of approximately 6,000 ever-married women.

Table A3: Summary Statistics for Men: Demographic Health Survey (DHS)

	All		Outside a Mill's CA		Within a Mill's CA	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Panel A: Main dependent variables						
Worked in the past 12 months	0.91	0.29	0.93	0.26	0.89	0.32
Worked for cash in the past 12 months	0.83	0.38	0.83	0.38	0.84	0.37
Panel B: Controls (Used as Controls)						
Christian	0.95	0.21	0.95	0.22	0.97	0.17
Age (Controlled via age group fixed effects)	37.94	9.88	37.78	9.86	38.40	9.97
Occupation: Agricultural self-employed	0.51	0.50	0.45	0.50	0.66	0.47
Education: No education	0.18	0.39	0.18	0.39	0.18	0.38
Education: Primary not completed	0.48	0.50	0.46	0.50	0.53	0.50
Education: Primary completed	0.20	0.40	0.21	0.41	0.19	0.39
Education: More than primary	0.14	0.34	0.15	0.36	0.10	0.30
Year of first marriage (Controlled via fixed effects)	1998.02	11.10	1998.78	11.00	1996.60	10.88
Panel C: Other Characteristics (Not Used as Controls)						
Marital status: Married	0.68	0.47	0.66	0.47	0.75	0.44
Marital status: Living together	0.29	0.45	0.31	0.46	0.22	0.41
Marital status: Divorced	0.02	0.13	0.02	0.13	0.02	0.13
Marital status: Not living together	0.02	0.13	0.02	0.14	0.02	0.13
Age at first marriage	24.24	4.72	24.20	4.77	24.49	4.54
Number of unions: One	0.86	0.35	0.87	0.34	0.86	0.35
Observations	9542		7245		1467	

Notes: This table presents summary statistics for men in my sample from the Demographic Health Survey (DHS). In each survey cycle, half of the households were randomly selected for the male survey, which is why the sample size for men is smaller compared to women. The majority of these men are the husbands of currently married women in the DHS sample, aligning closely with the female sample structure. Panel A reports the main dependent variables, while Panel B includes control variables used in the specification associated with the dataset. Age is controlled via age group fixed effects, measured in five-year age groups. Year of first marriage fixed effects (measured in years and collected for ever-married women) are included to account for shocks to the marriage market. I also include (men's) interview month and interview year fixed effects in the specification, where the results are robust to their exclusion. Panel C provides additional characteristics related to marital status and family dynamics (not used as controls). The sample is split by men living within or outside the catchment area of a mill (Mill's CA), which is defined as a 4 km radius buffer around a mill. Please refer to the details in Appendix A3.2 for the differences in the working-for-cash variable across women and men.

Table A4: Summary Statistics for Women: Integrated Household Living Conditions Survey (EICV)

	All		In a District without a Mill		In a District with a Mill	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Panel A: Main dependent variables						
Last daily earnings in Rwandan francs	172.15	1533.51	329.20	2729.82	104.30	358.64
Panel B: Characteristics (Used as Controls)						
Age (Controlled via age group fixed effects)	32.64	7.94	32.25	7.65	32.87	8.05
Occupation: At least one of the reported jobs is agricultural	0.96	0.19	0.96	0.20	0.97	0.18
Education: No education	0.18	0.39	0.13	0.34	0.21	0.41
Education: Incomplete primary	0.53	0.50	0.57	0.49	0.50	0.50
Education: Primary completed	0.18	0.39	0.16	0.36	0.20	0.40
Education: More than primary	0.10	0.30	0.14	0.34	0.09	0.28
Panel C: Other Characteristics (Not Used as Controls)						
Marital Status: Married	0.77	0.42	0.81	0.39	0.78	0.41
Marital status: Living together	0.18	0.38	0.15	0.36	0.17	0.37
Marital status: Divorce or separated	0.05	0.21	0.03	0.18	0.05	0.22
Observations	6363		1921		3686	

Notes: This table presents summary statistics for women in my sample from the Integrated Household Living Conditions Survey (EICV). Panel A reports the main dependent variable, last daily earnings in Rwandan francs, while Panel B includes control variables used in the specification associated with the dataset. Age is controlled via age group fixed effects, measured in five-year age groups. The year of marriage is not collected in EICV, and thus, year of marriage fixed effects are not included in specification 2. I also include (women's) interview month fixed effects in the specification, where the results are robust to their exclusion. Since the interview year and harvest year are the same for this dataset, only harvest year fixed effects are included. Panel C provides additional characteristics related to marital status (not used as controls). The sample is split by women living in a district with or without a mill. The average area for a district is 810 km^2. Zero values for the earnings variable are included in the sample.

Table A5: Summary Statistics for Men: Integrated Household Living Conditions Survey (EICV)

	All		In a District without a Mill		In a District with a Mill	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Panel A: Main dependent variables						
Last daily earnings in Rwandan francs	518.94	3084.56	915.32	4846.51	353.46	1979.38
Panel B: Characteristics (Used as Controls)						
Age (Controlled viaage group fixed effects)	37.75	10.26	36.74	10.14	38.24	10.33
Occupation: At least one of the reported jobs is agricultural	0.89	0.31	0.86	0.34	0.91	0.29
Education: No education	0.15	0.36	0.10	0.31	0.18	0.39
Education: Incomplete primary	0.52	0.50	0.56	0.50	0.49	0.50
Education: Primary completed	0.23	0.42	0.23	0.42	0.23	0.42
Education: More than primary	0.12	0.33	0.16	0.36	0.11	0.31
Panel C: Other Characteristics (Not used as Controls)						
Marital Status: Married	0.82	0.39	0.86	0.35	0.83	0.38
Marital status: Living together	0.17	0.38	0.14	0.35	0.16	0.37
Marital status: Divorce or separated	0.01	0.10	0.01	0.07	0.01	0.11
Observations	6685		2013		3886	

Notes: This table presents summary statistics for men in my sample from the Integrated Household Living Conditions Survey (EICV). Panel A reports the main dependent variable, last daily earnings in Rwandan francs, while Panel B includes control variables used in the specification associated with the dataset. Age is controlled via age group fixed effects, measured in five-year age groups. The year of marriage is not collected in EICV, and thus, year of marriage fixed effects are not included in specification 2. I also include (men's) interview month fixed effects in the specification, where the results are robust to their exclusion. Since the interview year and harvest year are the same for this dataset, only harvest year fixed effects are included. Panel C provides additional characteristics related to marital status (not used as controls). The sample is split by men living in a district with or without a mill. The average area for a district is 810 km^2. Zero values for the earnings variable are included in the sample.

Table A6: Summary Statistics for Hospitals: Hospital Management Information Sytem (HMIS)

	All		Outside a Mill's CA		Within a Mill's CA	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Panel A: Main dependent variables						
Has a GBV Patient: Women aged older than 18	0.86	0.35	0.88	0.32	0.81	0.39
Has a GBV Patient: Women aged 10-18 years old	0.89	0.32	0.90	0.30	0.85	0.35
Panel B: Breakdown of dependent variables in numbers						
Patients w. physical GBV symptoms: Women aged older than 18	5.50	7.78	6.18	8.82	4.10	4.67
Patients w. sexual GBV symptoms: Women aged older than 18	2.24	2.82	2.54	3.05	1.62	2.12
Patients w. physical GBV symptoms.: Women aged 10-18 years old	0.49	1.23	0.53	1.26	0.40	1.16
Patients w. sexual GBV symptoms: Women aged 10-18 years old	7.78	8.16	8.29	8.42	6.73	7.49
Observations	2557		1725		832	

Notes: This table presents summary statistics for gender-based violence (GBV) patients from the HMIS dataset. Panel A reports the proportion of hospitals with at least one GBV patient in a month, grouped by age group. Panel B provides the number of monthly hospitalizations for GBV, including both physical and sexual gender-based violence, for two age groups: women aged 18 and older (used as a proxy for domestic violence in the paper) and women aged 10-18 years old. "CA" stands for catchment area, which is defined as a 4 km radius buffer around mills.

A1.2 Additional Tables for Section 5: Results

Table A7: Dynamic Impact of a Mill Opening on Ever-Married Women Working and Working for Cash in the Past 12 Months

	Work-Related Variables Measured Over the Past 12 Months	
	(1) Working	(2) Working for Cash
Mill x Event Time -6- -3	0.03	-0.01
	(0.02)	(0.03)
Mill x Event Time -2	0.02	0.01
	(0.02)	(0.04)
Mill x Event Time -1	0.00	0.00
	(.)	(.)
Mill x Event Time 0-1	0.02	-0.03
	(0.02)	(0.03)
Mill x Event Time 2-3	0.01	0.06**
	(0.02)	(0.03)
Mill x Event Time 4-6	0.01	0.06**
	(0.02)	(0.03)
Observations	21052	18352
Dependent variable mean	0.88	0.37

Notes: This table reports coefficients from regressions of work-related outcomes on the interaction between residing in a mill's catchment area and event-year dummies, using equation (1). The dependent variables, drawn from the Demographic Health Surveys (DHS), measure whether a woman worked in the past 12 months (Column 1) and whether she worked for cash in the past 12 months (Column 2). The sample for working for cash is restricted to women who reported working in the past 12 months, as defined by the survey question, which explains the smaller number of observations in Column 2. The sample consists of ever-married women whose partner or most recent partner is alive. The event of interest is the opening of a mill. Event time is measured in years relative to the mill opening, with 0 indicating the year of the first mill's opening in the catchment area, negative values representing years before, and positive values representing years after. The omitted category is one year before the mill's opening. Event time years are grouped to ensure a balanced number of observations across years. All regressions control for individual characteristics, including women's religion, education, age-group fixed effects (measured in five-year age groups), survey month and year fixed effects, year of first marriage fixed effects, and sub-district fixed effects, and district-by-harvest year fixed effects. In Column 2, I control for agricultural occupation (Since I later show that agricultural occupation does not change with treatment, it is considered a predetermined characteristic/variable and can therefore be used as a control.). I cannot control for agricultural occupation when 'worked in the past 12 months' is the dependent variable, given that the occupation information is not collected for women who did not work in the past 12 months. This is the only dependent variable for which I cannot control for agricultural occupation in my DHS analysis. The catchment area is defined using a 4 km buffer around the mills. Standard errors are clustered at the sub-district level. *** p<0.01, ** p<0.05, * p<0.1.

Table A8: Dynamic Impact of a Mill Opening on Ever-Married Women Experiencing Self-Reported Domestic Violence in the Past 12 Months: Physical or Sexual, Physical Only, Sexual Only

	Experiencing Domestic Violence in the Past 12 Months		
	(1) Physical or Sexual	(2) Physical	(3) Sexual
Mill x Event Time -6- -2	-0.02	-0.00	-0.02
	(0.04)	(0.03)	(0.03)
Mill x Event Time -1	0.00	0.00	0.00
	(.)	(.)	(.)
Mill x Event Time 0-1	-0.13***	-0.09**	-0.07**
	(0.04)	(0.04)	(0.03)
Mill x Event Time 2-3	-0.10**	-0.09**	-0.05*
	(0.04)	(0.04)	(0.03)
Mill x Event Time 4-6	-0.07	-0.08**	-0.03
	(0.05)	(0.04)	(0.03)
Observations	5992	5991	5973
Dependent variable mean	0.32	0.29	0.11

Notes: This table reports coefficients from regressions of self-reported domestic violence outcomes on the interaction between residing in a mill's catchment area and event-year dummies, using equation (1). The dependent variables, drawn from the Demographic Health Surveys (DHS), measure whether a woman experienced physical or sexual violence (Column 1), physical violence only (Column 2), and sexual violence only (Column 3) in the past 12 months. The domestic violence module was administered to a randomly selected subsample of the full DHS sample, which explains the smaller number of observations compared to the work variables. Within this subsample, the analysis includes ever-married women whose partner or most recent partner is alive. Partnered women answer the questions for their current partners, while divorced or separated women answer them for their most recent partners. The event of interest is the opening of a mill. Event time is measured in years relative to the mill opening, with 0 indicating the year of the first mill's opening in the catchment area, negative values representing years before, and positive values representing years after. The omitted category is one year before the mill's opening. Event time years are grouped to ensure a balanced number of observations across years. All regressions control for individual characteristics, including women's religion, education, agricultural occupation, age-group fixed effects (measured in five-year age groups), survey month and year fixed effects, year of first marriage fixed effects, sub-district fixed effects, and district-by-harvest year fixed effects. The catchment area is defined using a 4 km buffer around the mills. Standard errors are clustered at the sub-district level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A9: Dynamic Impact of a Mill Opening on Ever-Married Women Working and Working for Cash in the Past 12 Months (Excluding Agricultural Occupation as a Control)

	Work-Related Variables Measured Over the Past 12 Months	
	(1) Working	(2) Working for Cash
Mill x Event Time -6- -3	0.03	-0.00
	(0.02)	(0.03)
Mill x Event Time -2	0.02	0.03
	(0.02)	(0.04)
Mill x Event Time -1	0.00	0.00
	(.)	(.)
Mill x Event Time 0-1	0.02	-0.02
	(0.02)	(0.03)
Mill x Event Time 2-3	0.01	0.07**
	(0.02)	(0.03)
Mill x Event Time 4-6	0.01	0.06**
	(0.02)	(0.03)
Observations	21052	20884
Dependent variable mean	0.88	0.33

Notes: This table reports coefficients from regressions of work-related outcomes on the interaction between residing in a mill's catchment area and event-year dummies, using equation (1). The dependent variables, drawn from the Demographic Health Surveys (DHS), measure whether a woman worked in the past 12 months (Column 1) and whether she worked for cash in the past 12 months (Column 2). The sample for working for cash is restricted to women who reported working in the past 12 months, as defined by the survey question, which explains the smaller number of observations in Column 2. The sample consists of ever-married women whose partner or most recent partner is alive. The event of interest is the opening of a mill. Event time is measured in years relative to the mill opening, with 0 indicating the year of the first mill's opening in the catchment area, negative values representing years before, and positive values representing years after. The omitted category is one year before the mill's opening. Event time years are grouped to ensure a balanced number of observations across years. All regressions control for individual characteristics, including women's religion, education, age-group fixed effects (measured in five-year age groups), survey month and year fixed effects, year of first marriage fixed effects, and sub-district fixed effects, and district-by-harvest year fixed effects. This table reports the same specification as Table A7, but excludes agricultural occupation from both regressions. Non-working women are coded as 0 in the working for cash variable. The catchment area is defined using a 4 km buffer around the mills. Standard errors are clustered at the sub-district level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A10: Dynamic Impact of a Mill Opening on Ever-Married Women Experiencing Self-Reported Domestic Violence in the Past 12 Months: Physical or Sexual, Physical Only, Sexual Only (Excluding Agricultural Occupation as a Control)

	Experiencing Domestic Violence in the Past 12 Months		
	(1) Physical or Sexual	(2) Physical	(3) Sexual
Mill x Event Time -6- -2	-0.04	-0.02	-0.02
	(0.04)	(0.03)	(0.02)
Mill x Event Time -1	0.00	0.00	0.00
	(.)	(.)	(.)
Mill x Event Time 0-1	-0.09**	-0.06*	-0.07**
	(0.04)	(0.03)	(0.03)
Mill x Event Time 2-3	-0.06	-0.06	-0.03
	(0.04)	(0.04)	(0.03)
Mill x Event Time 4-6	-0.06	-0.07*	-0.03
	(0.04)	(0.04)	(0.03)
Observations	6955	6954	6933
Dependent variable mean	0.31	0.29	0.11

Notes: This table reports coefficients from regressions of self-reported domestic violence outcomes on the interaction between residing in a mill's catchment area and event-year dummies, using equation (1). The dependent variables, drawn from the Demographic Health Surveys (DHS), measure whether a woman experienced physical or sexual violence (Column 1), physical violence only (Column 2), and sexual violence only (Column 3) in the past 12 months. The domestic violence module was administered to a randomly selected subsample of the full DHS sample, which explains the smaller number of observations compared to the work variables. Within this subsample, the analysis includes ever-married women whose partner or most recent partner is alive. Partnered women answer the questions for their current partners, while divorced or separated women answer them for their most recent partners. The event of interest is the opening of a mill. Event time is measured in years relative to the mill opening, with 0 indicating the year of the first mill's opening in the catchment area, negative values representing years before, and positive values representing years after. The omitted category is one year before the mill's opening. Event time years are grouped to ensure a balanced number of observations across years. All regressions control for individual characteristics, including women's religion, education, age-group fixed effects (measured in five-year age groups), survey month and year fixed effects, year of first marriage fixed effects, sub-district fixed effects, and district-by-harvest year fixed effects. This table reports the same specification as Table A8, but excludes agricultural occupation from all regressions. The catchment area is defined using a 4 km buffer around the mills. Standard errors are clustered at the sub-district level. *** p<0.01, ** p<0.05, * p<0.1.

Table A11: Dynamic Impact of a Mill Opening on Ever Married Women's Occupation

	(1) Managers	(2) Sales	(3) Agricultural Self-Employed	(4) Agricultural Employee	(5) Manual Skilled	(6) Manual Unskilled
Mill x Event Time -6- -3	0.00	0.00	0.02	-0.01	0.00	-0.02
	(0.01)	(0.02)	(0.03)	(0.01)	(0.01)	(0.02)
Mill x Event Time -2	-0.01	0.01	0.02	0.01	-0.01	-0.01
	(0.01)	(0.02)	(0.03)	(0.01)	(0.01)	(0.02)
Mill x Event Time -1	0.00	0.00	0.00	0.00	0.00	0.00
	(.)	(.)	(.)	(.)	(.)	(.)
Mill x Event Time 0-1	0.01	-0.00	-0.04*	0.01	-0.01	0.02
	(0.01)	(0.01)	(0.02)	(0.01)	(0.01)	(0.01)
Mill x Event Time 2-3	0.00	-0.02*	-0.00	0.01	-0.01	0.02
	(0.01)	(0.01)	(0.02)	(0.01)	(0.01)	(0.01)
Mill x Event Time 4-6	-0.00	-0.02	-0.00	0.00	-0.00	0.02*
	(0.01)	(0.01)	(0.02)	(0.01)	(0.00)	(0.01)
Observations	18387	18387	18387	18387	18387	18387
Dependent variable mean	0.03	0.07	0.74	0.08	0.02	0.06

Notes: This table reports coefficients from regressions of occupational categories on the interaction between residing in a mill's catchment area and event-year dummies, using equation (1). The dependent variables, drawn from the Demographic Health Surveys (DHS), indicate whether a woman's occupation falls into one of six categories: Managers (Column 1), Sales (Column 2), Agricultural Self-Employed (Column 3), Agricultural Employee (Column 4), Skilled Manual Worker (Column 5), and Unskilled Manual Worker (Column 6). The occupation question is asked to women who reported working. The sample consists of ever-married women whose partner or most recent partner is alive. The event of interest is the opening of a mill. Event time is measured in years relative to the mill opening, with 0 indicating the year of the first mill's opening in the catchment area, negative values representing years before, and positive values representing years after. The omitted category is one year before the mill opening. Event time years are grouped to ensure a balanced number of observations across years. All regressions control for individual characteristics, including women's religion, education, age-group fixed effects (five-year age groups), survey month and year fixed effects, year of first marriage fixed effects, sub-district fixed effects, and district-by-harvest-year fixed effects. The catchment area is defined using a 4 km buffer around the mills. Standard errors are clustered at the sub-district level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A12: Dynamic Impact of a Mill Opening on Ever-Married Men's Working, Working for Cash in the Past 12 Months and Occupation (Mostly Husbands of Ever-Married Women)

	Work		Occupation					
	(1) Work	(2) Cash Work	(3) Managers	(4) Sales	(5) Agricultural Self-Employed	(6) Agricultural Employee	(7) Manual Skilled	(8) Manual Unskilled
Mill x Event Time -6- -2	-0.03	0.02	0.02	0.02	0.00	0.00	-0.04	0.01
	(0.03)	(0.04)	(0.02)	(0.01)	(0.04)	(0.02)	(0.02)	(0.02)
Mill x Event Time -1	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	(.)	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Mill x Event Time 0-1	-0.01	0.01	0.01	0.01	-0.01	0.03	-0.04	0.01
	(0.04)	(0.03)	(0.02)	(0.02)	(0.04)	(0.02)	(0.03)	(0.02)
Mill x Event Time 2-3	-0.03	0.02	-0.01	0.02	-0.06*	0.02	-0.01	0.04
	(0.03)	(0.04)	(0.02)	(0.02)	(0.03)	(0.02)	(0.03)	(0.03)
Mill x Event Time 4-6	0.03	0.02	0.01	0.00	0.02	-0.01	-0.04*	0.02
	(0.03)	(0.03)	(0.01)	(0.01)	(0.03)	(0.02)	(0.02)	(0.02)
Observations	9525	8615	8622	8622	8622	8622	8622	8622
Dependent variable mean	0.91	0.83	0.05	0.05	0.51	0.13	0.13	0.10

Notes: This table reports coefficients from regressions of work-related outcomes on the interaction between residing in a mill's catchment area and event-year dummies, using equation (1). The dependent variables, drawn from the Demographic Health Surveys (DHS), measure whether a man worked in the past 12 months (Column 1) and whether he worked for cash in the past 12 months (Column 2). The sample for working for cash is restricted to men who reported working in the past 12 months, which explains the smaller number of observations in Column 2. From Columns 3-8, the table reports coefficients from regressions of occupational categories on the interaction between residing in a mill's catchment area and event-year dummies, using equation (1). The dependent variables indicate whether a man's occupation falls into one of six categories: Managers (Column 3), Sales (Column 4), Agricultural Self-Employed (Column 5), Agricultural Employee (Column 6), Skilled Manual Worker (Column 7), and Unskilled Manual Worker (Column 8). The occupation question is asked to men who reported working, which explains the smaller number of observations in Column 3-8 compared to Column 1. The sample consists of ever-married men whose partner or most recent partner is alive. The majority of these men are the husbands of currently married women in the DHS sample, aligning closely with the female sample structure. The event of interest is the opening of a mill. Event time is measured in years relative to the mill opening, with 0 indicating the year of the first mill's opening in the catchment area, negative values representing years before, and positive values representing years after. The omitted category is one year before the mill's opening. Event time years are grouped to ensure a balanced number of observations across years. All regressions control for individual characteristics, including men's religion, education, age-group fixed effects (measured in five-year age groups), survey month and year fixed effects, year of first marriage fixed effects, sub-district fixed effects, and district-by-harvest-year fixed effects. In Column 2, I additionally control for agricultural occupation, as this variable is only collected for men who report working. Since I show that agricultural occupation does not change with treatment, it can be considered a predetermined characteristic/variable in this specification and can therefore be used as a control. The catchment area is defined using a 4 km buffer around the mills. Standard errors are clustered at the sub-district level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A13: Dynamic Impact of a Mill Opening on Ever-Married Women Experiencing Self-Reported Domestic Violence in the Past 12 Months: Emotional Only

	Experiencing Domestic Violence in the Past 12 Months
	(1)
	Emotional
Mill x Event Time -6- -2	-0.03
	(0.04)
Mill x Event Time -1	0.00
	(.)
Mill x Event Time 0-1	-0.02
	(0.08)
Mill x Event Time 2-3	-0.01
	(0.04)
Mill x Event Time 4-6	0.00
	(0.05)
Observations	3468
Dependent variable mean	0.09

Notes: This table reports coefficients from regressions of self-reported emotional domestic violence on the interaction between residing in a mill's catchment area and event-year dummies. The dependent variable, drawn from the Demographic Health Surveys (DHS), indicates whether a woman experienced emotional domestic violence in the past 12 months. The sample consists of ever-married women whose partner or most recent partner is alive. Partnered women answer the questions for their current partners, while divorced or separated women answer them for their most recent partners. Unlike physical and sexual violence, emotional violence was not collected in the 2010 DHS round, which results in a smaller sample size compared to other domestic violence outcomes. The event of interest is the opening of a mill. Event time is measured in years relative to the mill opening, with 0 indicating the year of the first mill's opening in the catchment area, negative values representing years before, and positive values representing years after. The omitted category is one year before the mill opening. Event time years are grouped to ensure a balanced number of observations across years and match the grouping for household decision-making results. The regression control for individual characteristics, including women's religion, education, agricultural occupation, age-group fixed effects (five-year age groups), survey month and year fixed effects, year of first marriage fixed effects, sub-district fixed effects, and district-by-harvest-year fixed effects. The catchment area is defined using a 4 km buffer around the mills. Standard errors are clustered at the sub-district level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A14: Dynamic Impact of a Mill Opening on Couple's Last Daily Earnings (Wife + Husband) and the Wife's Share of Couple's Last Daily Earnings

	(1) Couple's Last Daily Earnings (Wife + Husband)	(2) Wife's Share of Couple's Last Daily Earnings
Mill in District x Event Time -7- -2	0.09	-0.03
	(0.18)	(0.02)
Mill in District x Event Time -1	0.00	0.00
	(.)	(.)
Mill in District x Event Time 0-1	-0.07	0.02
	(0.14)	(0.01)
Mill in District x Event Time 2-3	-0.33*	0.01
	(0.18)	(0.01)
Mill in District x Event Time 4-5	0.37***	0.03***
	(0.12)	(0.01)
Observations	5336	5336
Dependent variable mean	593.49	0.07

Notes: This table reports coefficients from the regression of couple's last daily earnings on the interaction between residing in a district with a mill and event-year dummies, using equation (2). The dependent variables are constructed from the Integrated Household Living Conditions Survey (EICV). Column 1 presents the results for the couple's last daily earnings, defined as the sum of the wife's and husband's last daily earnings, and is estimated using a Poisson pseudo-likelihood regression. Column 2 presents results for the wife's share of the couple's last daily earnings, defined as the ratio of the wife's last daily earnings to total couple earnings, and is estimated using a linear regression model. The event of interest is the opening of a mill in a district. I define spatial variation by whether couples reside in a district with a mill, not the catchment area, given that the data is not geocoded. A district encompasses a larger geographical area than both a catchment area and a sub-district. Following the opening of the first mill, the number of mills in a district increases over time, raising the likelihood that an individual is exposed to a mill. As expected, the mill's effects emerge a few years post-opening rather than immediately in year zero. Event time is measured in years, with 0 indicating the year of the first mill's opening in the district, negative values representing years before, and positive values representing years after. The omitted category is one year before the mill's opening. Event time years are grouped to ensure a balanced number of observations across years. The sample consists of linked wife-husband pairs. The regression controls for individual characteristics, including the wife's and husband's education, agricultural occupation, age-group fixed effects (measured in five-year age groups), survey month fixed effects, district fixed effects, and harvest year fixed effects. Standard errors are clustered at the district level. The earnings sample includes linked wife-husband pairs surveyed within 7 years before and 5 years after a mill opening, maintaining alignment with the DHS analysis event window. The sample excludes couples residing in Kigali (the capital) due to its distinct, highly urbanized labor market with earnings levels differing significantly from the rest of the country. Including Kigali yields similar results. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

This figure plots the coefficients from the regression of couple's last daily earnings on the interaction between residing in a district with a mill and event-year dummies, using equation (2).The dependent variables, drawn from the Integrated Household Living Conditions Survey (EICV), are estimated using a Poisson pseudo-likelihood regression. The dependent variables are constructed from the Integrated Household Living Conditions Survey (EICV). The left figure examines the couple's last daily earnings, defined as the sum of the wife's and husband's last daily earnings, and is estimated using a Poisson pseudo-likelihood regression. The right figure examines the wife's share of the couple's last daily earnings, defined as the ratio of the wife's last daily earnings to total couple earnings, and is estimated using a linear regression model. The event of interest is the opening of a mill, marked by the purple dashed line. Event time is measured in years, with 0 indicating the year of the mill's opening, negative values representing years before, and positive values representing years after. Event time years are grouped to ensure a balanced number of observations across years. The sample consists of linked wife-husband pairs. The y-axis shows the change in dependent variables relative to one year before the mill's opening (the omitted category). Standard errors are clustered at the district level, and 95% confidence intervals are displayed. The sample excludes individuals residing in Kigali, the capital, due to its highly urbanized labor market, where earnings levels differ significantly from the rest of the country. Including Kigali yields similar results.

Table A15: Dynamic Impact of a Mill Opening in a District on Ever Married Women's Last Daily Earnings

	(1) Woman: Last Daily Earnings
Mill in District x Event Time -7- -2	0.17 (0.16)
Mill in District x Event Time -1	0.00 (.)
Mill in District x Event Time 0-1	-0.41 (0.28)
Mill in District x Event Time 2-3	-0.07 (0.17)
Mill in District x Event Time 4-5	0.28** (0.12)
Observations	6211
Dependent variable mean	174.55

Notes: This table reports coefficients from the regression of women's last daily earnings on the interaction between residing in a district with a mill and event-year dummies, using equation (2). The dependent variable, drawn from the Integrated Household Living Conditions Survey (EICV), is estimated using a Poisson pseudo-likelihood regression. The event of interest is the opening of a mill in a district. I define spatial variation by whether individuals reside in a district with a mill, not the catchment area, given that the data is not geocoded. A district encompasses a larger geographical area than both a catchment area and a sub-district. Following the opening of the first mill, the number of mills in a district increases over time, raising the likelihood that an individual is exposed to a mill. As expected, the mill's effects emerge a few years post-opening rather than immediately in year zero. Event time is measured in years, with 0 indicating the year of the first mill's opening in the district, negative values representing years before, and positive values representing years after. The omitted category is one year before the mill's opening. Event time years are grouped to ensure a balanced number of observations across years. The sample consists of ever-married women whose partner or most recent partner is alive. The regression controls for individual characteristics, including women's education, agricultural occupation, age-group fixed effects (measured in five-year age groups), survey month fixed effects, district fixed effects, and harvest year fixed effects. Standard errors are clustered at the district level. The earnings sample includes individuals surveyed within 7 years before and 5 years after a mill opening, maintaining alignment with the DHS analysis event window. The sample excludes individuals residing in Kigali (the capital) due to its distinct, highly urbanized labor market with earnings levels differing significantly from the rest of the country. Including Kigali yields similar results, although, the confidence intervals for the insignificant estimates in some event years are larger. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A16: Dynamic Impact of a Mill Opening in a District on Ever-Married Men's Last Daily Earnings

	(1) Man: Last Daily Earnings
Mill in District x Event Time -7- -2	0.36 (0.63)
Mill in District x Event Time -1	0.00 (.)
Mill in District x Event Time 0-1	0.16 (0.31)
Mill in District x Event Time 2-3	-0.31 (0.42)
Mill in District x Event Time 4-5	0.67 (0.45)
Observations	4902
Dependent variable mean	473.92

Notes: This table reports coefficients from the regression of men's last daily earnings on the interaction between residing in a district with a mill and event-year dummies, using equation (2). The dependent variable, drawn from the Integrated Household Living Conditions Survey (EICV), is estimated using a Poisson pseudo-likelihood regression. The event of interest is the opening of a mill in a district. I define spatial variation by whether individuals reside in a district with a mill, not the catchment area, given that the data is not geocoded. A district encompasses a larger geographical area than both a catchment area and a sub-district. Following the opening of the first mill, the number of mills in a district increases over time, raising the likelihood that an individual is exposed to a mill. As expected, the mill's effects emerge a few years post-opening rather than immediately in year zero. Event time is measured in years, with 0 indicating the year of the first mill's opening in the district, negative values representing years before, and positive values representing years after. The omitted category is one year before the mill's opening. Event time years are grouped to ensure a balanced number of observations across years. The sample consists of ever-married men whose partner or most recent partner is alive, mostly the husbands of married women in the sample. The regression controls for individual characteristics, including men's education, agricultural occupation, age-group fixed effects (measured in five-year age groups), survey month fixed effects, district fixed effects, and harvest year fixed effects. Standard errors are clustered at the district level. The earnings sample includes individuals surveyed within 7 years before and 5 years after a mill opening, maintaining alignment with the DHS analysis event window. The sample excludes individuals residing in Kigali (the capital) due to its distinct, highly urbanized labor market with earnings levels differing significantly from the rest of the country. Including Kigali yields similar results, although, the confidence intervals for the insignificant estimates in some event years are larger. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A17: Dynamic Impact of the Start of Mill Operations on Women's Hospitalizations for Gender Based Violence

	Gender-Based Violence Hospitalizations: Women	
	(1) Older than 18 Years Old	(2) 18 Years Old and Younger
Hospital in Mill Catchment x Month 1	0.00	-0.00
	(0.05)	(0.04)
Hospital in Mill Catchment x Month 2	0.00	0.00
	(.)	(.)
Hospital in Mill Catchment x Month 3	-0.01	-0.06
	(0.04)	(0.06)
Hospital in Mill Catchment x Month 4	0.00	0.03
	(0.05)	(0.04)
Hospital in Mill Catchment x Month 5	-0.03	-0.01
	(0.05)	(0.06)
Hospital in Mill Catchment x Month 6	-0.11**	-0.06
	(0.04)	(0.05)
Hospital in Mill Catchment x Month 7	-0.14***	-0.02
	(0.04)	(0.05)
Hospital in Mill Catchment x Month 8	0.01	-0.02
	(0.04)	(0.06)
Hospital in Mill Catchment x Month 9	-0.00	0.05
	(0.05)	(0.07)
Hospital in Mill Catchment x Month 10	0.02	0.02
	(0.04)	(0.06)
Hospital in Mill Catchment x Month 11	-0.04	-0.01
	(0.08)	(0.06)
Hospital in Mill Catchment x Month 12	-0.02	0.05
	(0.06)	(0.07)
Observations	2538	2553
Dependent variable mean	0.86	0.89

Notes: This table reports coefficients from the regression of hospitalizations for (physical or sexual) gender-based violence for women on the interaction between a hospital being in a mill's catchment area and calendar-month dummies, using equation (3). The dependent variable, drawn from the Hospital Management Information System (HMIS), is a binary indicator for whether a hospital recorded any hospitalization due to gender-based violence for women. Column 1 presents results for women older than 18 (a proxy for domestic violence), while Column 2 presents results for women 18 years old and younger. The sample consists of 43 public hospitals in Rwanda that recorded monthly gender-based violence cases during 2012-2019, out of the 48 public hospitals in the country by 2019, representing the universe of such hospital records in the country. The event of interest is the start of the mill's operations, which begin in March and coincide with the start of the coffee harvest season. The harvest season peaks between May and July and is followed by the post-harvest period from August to December, when mill operations have ended. Calendar months are used, with February as the reference month. All regressions control for hospital fixed effects, month fixed effects, and province-by-harvest year fixed effects. The catchment area is defined using a 4 km buffer around the mills. Standard errors are clustered at the district level. Kabgayi District Hospital switched from being outside to within a mill's catchment area following a mill opening nearby in 2017. The results are robust to both correcting the treatment/control group status (the approach used here) and dropping this hospital from the sample. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A18: Dynamic Impact of the Start of Mill Operations on Women's Hospitalizations for Bone and Joint Disorders other than Fractures for Women Older than 18 (Non-Gender Based Violence Hospitalizations)

	Bone and Joint Disorders other than Fractures
	(1)
	Women Older than 18
Hospital in Mill Catchment x Month 1	0.11
	(0.15)
Hospital in Mill Catchment x Month 2	0.00
	(.)
Hospital in Mill Catchment x Month 3	0.06
	(0.09)
Hospital in Mill Catchment x Month 4	0.01
	(0.13)
Hospital in Mill Catchment x Month 5	0.08
	(0.12)
Hospital in Mill Catchment x Month 6	0.08
	(0.18)
Hospital in Mill Catchment x Month 7	0.03
	(0.16)
Hospital in Mill Catchment x Month 8	-0.12
	(0.16)
Hospital in Mill Catchment x Month 9	0.01
	(0.16)
Hospital in Mill Catchment x Month 10	0.13
	(0.18)
Hospital in Mill Catchment x Month 11	0.04
	(0.14)
Hospital in Mill Catchment x Month 12	0.09
	(0.16)
Observations	1523
Dependent variable mean	20.96

Notes: This table reports coefficients from the regression of hospitalizations for bone and joint disorders other than fractures (e.g., osteoarthritis, gout, rheumatoid arthritis, lupus, bursitis) for women older than 18 on the interaction between a hospital being in a mill's catchment area and calendar-month dummies, using equation (3). The dependent variable is drawn from the Hospital Management Information System (HMIS) (2012-2015) and defined by the Ministry of Health hospital classification as "Bone and joint disorders other than fractures". These cases are not recorded under the "Gender-Based Violence" section in the monthly hospital report forms. The variable is measured in logs. The sample consists of 43 public hospitals in Rwanda that recorded monthly gender-based violence cases during 2012-2019, out of the 48 public hospitals in the country by 2019. The event of interest is the start of the mill's operations, which begin in March and coincide with the start of the coffee harvest season. The harvest season peaks between May and July and is followed by the post-harvest period from August to December, when mill operations have ended. Calendar months are used, with February as the reference month. All regressions control for hospital fixed effects, month fixed effects, and province-by-harvest year fixed effects. The catchment area is defined using a 4 km buffer around the mills. Standard errors are clustered at the district level. *** p<0.01, ** p<0.05, * p<0.1.

Table A19: Dynamic Impact of the Start of Mill Operations on All Hospitalizations (Gender-Based and Non-Gender-Based Violence) for Women Older than 18

	All Hospitalizations: Women
	(1)
	Older than 18
Hospital in Mill Catchment x Month 1	-0.03
	(0.05)
Hospital in Mill Catchment x Month 2	0.00
	(.)
Hospital in Mill Catchment x Month 3	0.04
	(0.04)
Hospital in Mill Catchment x Month 4	0.02
	(0.04)
Hospital in Mill Catchment x Month 5	0.05
	(0.05)
Hospital in Mill Catchment x Month 6	-0.04
	(0.04)
Hospital in Mill Catchment x Month 7	0.02
	(0.06)
Hospital in Mill Catchment x Month 8	0.04
	(0.04)
Hospital in Mill Catchment x Month 9	0.08
	(0.05)
Hospital in Mill Catchment x Month 10	0.03
	(0.07)
Hospital in Mill Catchment x Month 11	-0.02
	(0.05)
Hospital in Mill Catchment x Month 12	0.07
	(0.07)
Observations	1216
Dependent variable mean	395.89

Notes: This table reports coefficients from the regression of all recorded hospitalizations (e.g., gender-based and non-gender based) for women older than 18 on the interaction between a hospital being in a mill's catchment area and calendar-month dummies, using equation (3). The dependent variable is drawn from the Hospital Management Information System (HMIS) and is recorded as the total number of hospitalizations in a given month for the women older than 18. The variable is measured in logs. The sample consists of 43 public hospitals in Rwanda that recorded monthly gender-based violence cases during 2012-2019, out of the 48 public hospitals in the country by 2019. The event of interest is the start of the mill's operations, which begin in March and coincide with the start of the coffee harvest season. The harvest season peaks between May and July and is followed by the post-harvest period from August to December, when mill operations have ended. Calendar months are used, with February as the reference month. All regressions control for hospital fixed effects, month fixed effects, and province-by-harvest year fixed effects. The catchment area is defined using a 4 km buffer around the mills. Standard errors are clustered at the district level. The analysis includes the years 2016–2019 instead of 2012–2019 due to data availability and sample size considerations. While the primary HMIS sample excludes 2016 and 2017 (drought years), I drop only 2017 here (the more severe drought year) to ensure sufficient statistical power. This robustness check examines whether women were systematically less likely to seek hospital care during the harvest months due to an increased opportunity cost of going to the hospital. In Rwanda, healthcare is accessible through a community-based health insurance system, making a public hospital care nearly free for most citizens. Additionally, drought conditions could reduce work opportunities rather than increase the opportunity cost of seeking medical care, potentially making women more likely to visit hospitals (e.g., non-gender-based violence cases). Therefore, finding no decline in hospital visits during the harvest months serves as a meaningful validity check, suggesting that women do not systematically avoid hospitals during this period. Kabgayi District Hospital switched from being outside to within a mill's catchment area following a mill opening nearby in 2017. The results are robust to both correcting the treatment/control group status (the approach used here) and dropping this hospital from the sample. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A20: Dynamic Impact of the Start of Mill Operations on Non-Gender-Based Violence Injuries: Falls

	Non-GBV Injuries: Falls
	(1)
	Women Older than 19
Mill x Month 1	-0.06
	(0.04)
Mill x Month 2	0.00
	(.)
Mill x Month 3	0.01
	(0.04)
Mill x Month 4	-0.02
	(0.03)
Mill x Month 5	-0.02
	(0.05)
Mill x Month 6	-0.06
	(0.04)
Mill x Month 7	-0.04
	(0.05)
Mill x Month 8	-0.05
	(0.06)
Mill x Month 9	-0.02
	(0.04)
Mill x Month 10	0.03
	(0.06)
Mill x Month 11	-0.04
	(0.05)
Mill x Month 12	0.02
	(0.05)
Observations	1198
Dependent variable mean	0.39

Notes: This table reports coefficients from the regression of hospitalizations for non-gender-based violence injuries for women older than 19 on the interaction between a hospital being in a mill's catchment area and calendar-month dummies, using equation (3). The dependent variable is drawn from the Hospital Management Information System (HMIS) and is not recorded under the "Gender-Based Violence" section in the monthly hospital report forms; instead, these cases are recorded under the "Injuries" category. The dependent variable is a binary indicator for whether a hospital recorded any hospitalization due to injuries due to falls for women. The sample consists of 43 public hospitals in Rwanda that recorded monthly gender-based violence cases during 2012-2019, out of the 48 public hospitals in the country by 2019. The event of interest is the start of the mill's operations, which begin in March and coincide with the start of the coffee harvest season. The harvest season peaks between May and July and is followed by the post-harvest period from August to December, when mill operations have ended. Calendar months are used, with February as the reference month. All regressions control for hospital fixed effects, month fixed effects, and province-by-harvest year fixed effects. The catchment area is defined using a 4 km buffer around the mills. Standard errors are clustered at the district level. The analysis includes the years 2016–2019 instead of 2012–2019 due to data availability and sample size considerations. While the primary HMIS sample excludes 2016 and 2017 (drought years), I drop only 2017 here (the more severe drought year) to ensure sufficient statistical power. This robustness check examines whether women were systematically less likely to seek hospital care during the harvest months due to an increased opportunity cost of going to the hospital. In Rwanda, healthcare is accessible through a community-based health insurance system, making a public hospital care nearly free for most citizens. Additionally, drought conditions could reduce work opportunities rather than increase the opportunity cost of seeking medical care, potentially making women more likely to visit hospitals (e.g., non-gender-based violence cases). Therefore, finding no decline in hospital visits during the harvest months serves as a meaningful validity check, suggesting that women do not systematically avoid hospitals during this period. Kabgayi District Hospital switched from being outside to within a mill's catchment area following a mill opening nearby in 2017. The results are robust to both correcting the treatment/control group status (the approach used here) and dropping this hospital from the sample. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A21: Dynamic Impact of a Mill Opening on Ever-Married Women's Help Seeking Behavior Following Domestic Violence (Question Asked to Women Who Reported Ever Experiencing Domestic Violence)

	(1) Ever Sought Help for Domestic Violence
Mill x Event Time -6- -2	0.11
	(0.08)
Mill x Event Time -1	0.00
	(.)
Mill x Event Time 0-1	-0.03
	(0.07)
Mill x Event Time 2-3	-0.02
	(0.10)
Mill x Event Time 4-6	-0.02
	(0.06)
Observations	3204
Dependent variable mean	0.44

Notes: This table reports coefficients from regressions of self-reported help seeking behavior on domestic violence on the interaction between residing in a mill's catchment area and event-year dummies, using equation (1). The dependent variable, drawn from the Demographic Health Surveys (DHS), measures whether a woman who experienced physical or sexual violence ever sought help from anyone (own family, husband's family, current or former partner, friend, neighbor, religious leader, doctor or medical personnel, police, lawyer or a social service organization). The help-seeking question is asked only to women who report experiencing physical or sexual violence. The domestic violence module was administered to a randomly selected subsample of the full DHS sample, which explains the smaller number of observations compared to the work variables. Within this subsample, the analysis includes ever-married women whose partner or most recent partner is alive. Partnered women answer the questions for their current partners, while divorced or separated women answer them for their most recent partners. The event of interest is the opening of a mill. Event time is measured in years relative to the mill opening, with 0 indicating the year of the first mill's opening in the catchment area, negative values representing years before, and positive values representing years after. The omitted category is one year before the mill's opening. Event time years are grouped to ensure a balanced number of observations across years. All regressions control for individual characteristics, including women's religion, education, agricultural occupation, age-group fixed effects (measured in five-year age groups), survey month and year fixed effects, year of first marriage fixed effects, sub-district fixed effects, and district-by-harvest year fixed effects. The catchment area is defined using a 4 km buffer around the mills. Standard errors are clustered at the sub-district level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

A1.3 Additional Tables for Section 6: Mechanisms

Table A22: Dynamic Impact of a Mill Opening on Married Women's Participation in Household Decision Making (Alone or Jointly with the Husband): All, Financial and Non-Financial Decisions

	Household Decisions Index		
	(1) All Decisions	(2) Financial Decisions: Large Household Purchases	(3) Non-Financial Decisions: Family Visit, Own Health, Contraception
Mill x Event Time -7- -2	0.03	0.00	0.03
	(0.03)	(0.04)	(0.04)
Mill x Event Time -1	0.00	0.00	0.00
	(.)	(.)	(.)
Mill x Event Time 0-1	0.12***	0.16***	0.10**
	(0.04)	(0.06)	(0.04)
Mill x Event Time 2-4	0.01	0.01	0.01
	(0.04)	(0.05)	(0.04)
Mill x Event Time 5-7	0.02	-0.03	0.04
	(0.04)	(0.06)	(0.04)
Observations	16525	16509	16524
Dependent variable mean	0.23	0.22	0.23

Notes: This table reports coefficients from regressions of household decision-making outcomes on the interaction between residing in a mill's catchment area and event-year dummies, using equation (1). The dependent variables are drawn from the Demographic Health Surveys (DHS) and are measured as standardized indices. A higher index value indicates that the wife is more likely to make decisions either alone or jointly with her husband, rather than her husband or someone else in the family usually making the decisions for her. Column 1 presents an index of all household decisions, Column 2 focuses on decisions regarding large household purchases (financial decisions), and Column 3 includes decisions about family visits, own healthcare, and contraception (non-financial decisions). Since this question is collected only among couples, the sample consists of married women. Event time is measured in years relative to the mill opening, with 0 indicating the year of the first mill's opening in the catchment area, negative values representing years before, and positive values representing years after. The omitted category is one year before the mill's opening. Event time years are grouped to ensure a balanced number of observations across years. All regressions control for individual characteristics, including women's religion, education, agricultural occupation, age-group fixed effects (measured in five-year age groups), survey month and year fixed effects, year of first marriage fixed effects, sub-district fixed effects, and district-by-harvest year fixed effects. The catchment area is defined using a 4 km buffer around the mills. Standard errors are clustered at the sub-district level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A23: Dynamic Impact of a Mill Opening on Married Women's Participation in Household Decision Making (Alone or Jointly with the Husband; Contraception Decision: Joint Only) All Decisions

	Household Decisions Index
	(1)
	All Decisions (Contraception is a Joint Decision)
Mill x Event Time -7- -2	0.02
	(0.03)
Mill x Event Time -1	0.00
	(.)
Mill x Event Time 0-1	0.11**
	(0.04)
Mill x Event Time 2-4	0.01
	(0.04)
Mill x Event Time 5-7	0.02
	(0.04)
Observations	16525
Dependent variable mean	0.23

Notes: This table reports coefficients from regressions of household decision-making outcomes on the interaction between residing in a mill's catchment area and event-year dummies, using equation (1). The dependent variable is drawn from the Demographic Health Surveys (DHS) and is measured as a standardized index. A higher index value indicates that the wife is more likely to make decisions either alone or jointly with her husband, rather than her husband or someone else in the family usually making the decisions for her. For large household purchases, family visits, and own health decisions, the variable is constructed as whether the wife is more likely to make decisions alone or jointly with her husband. For contraception decisions, the variable is constructed as whether the wife is more likely to make decisions jointly with her husband. Column 1 presents an index of all household decisions (large household purchases, family visit, own health, and contraception). Since this question is collected only among couples, the sample consists of married women (cohabitating couples included). Event time is measured in years relative to the mill opening, with 0 indicating the year of the first mill's opening in the catchment area, negative values representing years before, and positive values representing years after. The omitted category is one year before the mill's opening. Event time years are grouped to ensure a balanced number of observations across years. All regressions control for individual characteristics, including women's religion, education, agricultural occupation, age-group fixed effects (measured in five-year age groups), survey month and year fixed effects, year of first marriage fixed effects, sub-district fixed effects, and district-by-harvest year fixed effects. The catchment area is defined using a 4 km buffer around the mills. Standard errors are clustered at the sub-district level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A24: Dynamic Impact of a Mill Opening on Ever-Married Women's Current Marital Status: Divorced or Separated

	(1) Current Marital Status: Divorced or Separated
Mill x Event Time -7- -2	0.01
	(0.02)
Mill x Event Time -1	0.00
	(.)
Mill x Event Time 0-1	0.02
	(0.02)
Mill x Event Time 2-4	0.01
	(0.01)
Mill x Event Time 5-7	0.00
	(0.02)
Observations	19014
Dependent variable mean	0.13

Notes: This table reports coefficients from regressions of women's current marital status being divorced or separated on the interaction between residing in a mill's catchment area and event-year dummies, using equation (1). The dependent variable, drawn from the Demographic Health Surveys (DHS), is a binary indicator for whether a woman is currently divorced or separated at the time of the survey. The sample consists of ever-married women whose partner or most recent partner is alive. Event time is measured in years relative to the mill opening, with 0 indicating the year of the first mill's opening in the catchment area, negative values representing years before, and positive values representing years after. The omitted category is one year before the mill opening. Event time years are grouped to ensure a balanced number of observations across years. All regressions control for individual characteristics, including women's religion, education, agricultural occupation, age-group fixed effects (five-year age groups), survey month and year fixed effects, year of first marriage fixed effects, sub-district fixed effects, and district-by-harvest-year fixed effects. The catchment area is defined using a 4 km buffer around the mills. Standard errors are clustered at the sub-district level. I used the minimum and maximum event years –7 and 7 when using the divorce outcome variable, as in the household decision-making analysis, given that I discuss these two variables in parallel in Section 6.1. I used these specific minimum and maximum event years (instead of –6 and 6, as in the main results) because the household decision-making variables are defined for married women only, resulting in a smaller sample size compared to the sample in the main results. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A25: Dynamic Impact of a Mill Opening on Ever Married Women's Working for Cash and Self-Reported Domestic Violence in the Past 12 Months: Women with Different Occupations from Their Current or Most Recent Husband

	(1) Working for Cash	(2) Physical or Sexual Domestic Violence
Mill x Event Time -13- -2	0.04	-0.01
	(0.04)	(0.08)
Mill x Event Time -1	0.00	0.00
	(.)	(.)
Mill x Event Time 0-2	-0.06	0.01
	(0.04)	(0.09)
Mill x Event Time 3-5	0.07*	-0.13*
	(0.04)	(0.08)
Observations	6728	1953
Dependent variable mean	0.28	0.30

Notes: This table reports coefficients from regressions of women's working for cash and self-reported physical or sexual domestic violence in the past 12 months on the interaction between residing in a mill's catchment area and event-year dummies, using equation (1). The dependent variables are drawn from the Demographic Health Surveys (DHS). The domestic violence module was administered to a randomly selected subsample of the full DHS sample, which explains the smaller sample size in Column 2 compared to Column 1. The sample consists of ever-married women whose partner or most recent partner is alive and who do not share the same occupation as their partner (or most recent partner). In the survey, women report their own occupation, while for their partner's occupation, they report the occupation of their current spouse if married or their most recent partner if divorced or separated. Partnered women answer the domestic violence questions for their current partners, while divorced or separated women answer them for their most recent partners. Self-employed women working in agriculture (e.g., smallholder farmers), whose husbands have skilled or unskilled manual jobs (e.g., truck drivers, plumbers, refuse workers, laborers in mining, construction, and manufacturing), constitute the majority of this subsample. The event of interest is the opening of a mill. Event time is measured in years relative to the mill opening, with 0 indicating the year of the first mill's opening in the catchment area, negative values representing years before, and positive values representing years after. The omitted category is one year before the mill's opening. Event time years are grouped to ensure a balanced number of observations across years and to maintain a sufficient sample size, given the smaller subsample. All regressions control for individual characteristics, including women's religion, education, agricultural occupation, age-group fixed effects (measured in five-year age groups), survey month and year fixed effects, year of first marriage fixed effects, sub-district fixed effects, and district-by-harvest year fixed effects. The catchment area is defined using a 4 km buffer around the mills. Standard errors are clustered at the sub-district level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A26: Dynamic Impact of a Mill Opening on Husbands' Last Daily Earnings: Couples with Different Occupations

	(1) Husband: Last Daily Earnings
Mill in District x Event Time -7- -2	0.11
	(0.14)
Mill in District x Event Time -1	0.00
	(.)
Mill in District x Event Time 0-3	0.16
	(0.11)
Mill in District x Event Time 4-6	0.12
	(0.22)
Mill in District x Event Time 7-10	0.16
	(0.30)
Observations	6572
Dependent variable mean	1657.61

Notes: This table reports coefficients from the regression of men's last daily earnings on the interaction between residing in a district with a mill and event-year dummies, using equation (2). The dependent variable, drawn from the Integrated Household Living Conditions Survey (EICV), is estimated using a Poisson pseudo-likelihood regression. The event of interest is the opening of a mill in a district. I define spatial variation by whether individuals reside in a district with a mill, not the catchment area, given that the data is not geocoded. Event time is measured in years, with 0 indicating the year of the first mill's opening in the district, negative values representing years before, and positive values representing years after. The omitted category is one year before the mill's opening. Event time years are grouped to ensure a balanced number of observations across years. The subsample consists of husbands of currently married women in the EICV sample who do not share the same agricultural occupation as their wives. Unfortunately, the EICV dataset does not collect occupation information for the most recent partner of a divorced or separated man. To ensure consistency with the rule I applied in the DHS (where I made sample restrictions conditional on knowing the occupation information of both partners and ex-partners), I first restrict the EICV sample to currently married men only. Among these married men, I drop any man if both he and his wife have an agricultural occupation. This process results in a small reduction in the subsample size, approximately 300 divorced or separated men whose occupation is non-agricultural, for whom I do not know their former partner's occupation. The results remain similar when I include them in the subsample. This exercise allows me to confirm that, among the subsample of couples where most women work in agriculture and their partners do not, household earnings do not increase through the husband, meaning there is no change in husbands' earnings. Since DHS results show a decrease in domestic violence within this subsample, an increase in household income is also a plausible mechanism behind the results. This table corresponds to the bottom panel in Figure 12. The regression controls for individual characteristics, including men's education, age-group fixed effects (measured in five-year age groups), survey month fixed effects, district fixed effects, and harvest year fixed effects. Kigali is not excluded, as the aim here is to create the same subsample as in the DHS analysis in Table A20, and the DHS subsample does not exclude it. In the upcoming triple-difference earnings specification, Kigali is excluded, and the earnings results for this group of men remain unchanged after a mill opening. There are more post-mill opening years in the earnings analysis to ensure a sufficient sample size. Standard errors are clustered at the district level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table A27: Dynamic Impact of a Mill Opening on the Men's Last Daily Earnings: The Wife Has an Agricultural Occupation and the Husband has a Non-Agricultural Occupation

	(1) Men: Last Daily Earnings
Mill in District x Event Time -7- -2 x Agri Wife and Non-Agri Husband	0.15
	(0.25)
Mill in District x Event Time -1 x Agri Wife and Non-Agri Husband	0.00
	(.)
Mill in District x Event Time 0-2 x Agri Wife and Non-Agri Husband	0.12
	(0.29)
Mill in District x Event Time 3-5 x Agri Wife and Non-Agri Husband	-0.10
	(0.34)
Observations	6416
Dependent variable mean	533.61

Notes: This table reports coefficients from the regression of men's last daily earnings on the interaction between residing in a district with a mill, event-year dummies, and a third interaction for being in a couple where the wife works in agriculture and the husband works in a non-agricultural occupation. This regression builds on equation (2) by incorporating a third interaction term. The dependent variable, drawn from the Integrated Household Living Conditions Survey (EICV), is estimated using a Poisson pseudo-likelihood regression. The event of interest is the opening of a mill in a district. I define spatial variation by whether individuals reside in a district with a mill, not the catchment area, given that the data is not geocoded. Event time is measured in years, with 0 indicating the year of the first mill's opening in the district, negative values representing years before, and positive values representing years after. The omitted category is one year before the mill's opening. Event time years are grouped to ensure a balanced number of observations across years and to have a sufficient amount of agricultural wife and non-agricultural husband couples/observations in each event year. The sample consists of ever-married men whose partner or most recent partner is alive, most of whom are the husbands of married women in the sample. The majority of couples in the DHS subsample who do not share occupations are those where the wife has an agricultural occupation while the husband does not. This check confirms that, particularly for these couples, there is no change in the husband's earnings. This specification introduces a third interaction between mill exposure, event time, and an indicator for couples in which the wife is employed in agriculture while the husband works in a non-agricultural occupation. This robustness check further strengthens the interpretation that the reduction in domestic violence within this subgroup is not driven by an increase in the husband's earnings, reinforcing the plausibility of the income effect through the wife's earnings as a mechanism behind the observed reduction in violence. The regression controls for individual characteristics, including men's education, age-group fixed effects (measured in five-year age groups), survey month fixed effects, district fixed effects, and harvest year fixed effects. Standard errors are clustered at the district level. The sample excludes individuals residing in Kigali (the capital) due to its distinct, highly urbanized labor market with earnings levels differing significantly from the rest of the country. Including Kigali yields similar results. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A28: Dynamic Impact of the Start of Mill Operations on a Household's Monthly Food Consumption Expenditure

	(1) Monthly Food Consumption Expenditure
Mill in District x Month 1	0.03
	(0.04)
Mill in District x Month 2	0.00
	(.)
Mill in District x Month 3	0.01
	(0.02)
Mill in District x Month 4	0.04
	(0.04)
Mill in District x Month 5	-0.06*
	(0.03)
Mill in District x Month 6	0.03
	(0.02)
Mill in District x Month 7	0.80***
	(0.07)
Mill in District x Month 8	0.28***
	(0.09)
Mill in District x Month 9	-0.02
	(0.07)
Observations	58661
Dependent variable mean	7.23

Notes: This table reports coefficients from the regression of monthly household food consumption expenditure on the interaction between a district having a mill and calendar-month dummies, using an adjusted version of equation (3). The dependent variable, drawn from the Integrated Household Living Conditions Survey (EICV), is measured in logs. Since the time variation is at the monthly level rather than yearly, the increase in food consumption expenditure likely reflects higher food consumption rather than changes in food prices. The sample consists of ever-married women (whose partner or most recent partner is alive) with an agricultural occupation (the majority of the full sample), as the variation in the data is tied to a harvest season. The results are consistent even without this sample restriction, with effects remaining statistically significant for July and August. The variable is collected after the 2005 data cycle (2011, 2014, 2017 cycles), which is desirable since the gender-based violence hospitalization results—used alongside these expenditure results to analyze the incapacitation cost channel—are for 2012–2019. The event of interest is the start of mill operations, which coincides with the beginning of the coffee harvest season in March. The harvest season peaks between May and July and is followed by the post-harvest period from August to December, when mill operations have ended. Calendar months are used, with February as the reference month. The regression controls for individual characteristics, including women's education, age-group fixed effects (measured in five-year age groups), month fixed effects, district fixed effects, and harvest year fixed effects. Standard errors are clustered at the district level. The sample for the consumption expenditure variable includes approximately 14,500 households. The number of households in the consumption expenditure sample is much higher than in the earnings sample. The reason is as follows: Unlike the earnings variable, which is used in a DID event-study where the event of interest is a mill opening in a district, I use a different DID event-study for the food consumption expenditure variable (the same one used for hospitalizations), where the event of interest is the beginning of the harvest season (mill operations) within the year. In this mechanism check, the focus is on capturing monthly expenditures during the sample period for the hospitalization data (2012-2019), rather than the six years before and after a mill opening. Since this event-study specification does not restrict households to those who responded to the survey within 6 years after a mill opening (e.g., households in districts where a mill opened 7-10 years ago are included), the sample is larger here. The consumption expenditure results are robust to their exclusion. (The consumption expenditure results—showing higher consumption expenditure compared to pre-harvest levels, particularly in August when domestic violence hospitalizations revert to pre-harvest levels—are robust to excluding households in districts where a mill opened more than 6 years ago, ensuring the robustness of the evidence supporting the incapacitation channel.) *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A29: Effect of the Irish Potato Harvest on Gender-Based Violence Hospitalizations for Women Older than 18 in Potato-Suitable Districts

	Gender-Based Violence Hospitalizations: Women
	(1)
	Older than 18
Month 1	0.02
	(0.04)
Month 2	0.00
	(.)
Month 3	0.02
	(0.04)
Month 4	-0.03
	(0.05)
Month 5	0.00
	(0.03)
Month 6	-0.01
	(0.03)
Month 7	0.02
	(0.03)
Month 8	0.04
	(0.03)
Month 9	0.01
	(0.03)
Month 10	0.03
	(0.03)
Month 11	-0.00
	(0.04)
Month 12	0.01
	(0.04)
Observations	1588
Dependent variable mean	0.84

Notes: This table reports coefficients from the regression of gender-based violence hospitalizations for women older than 18 on calendar-month dummies in potato-suitable districts. The dependent variable, drawn from the Hospital Management Information System (HMIS), is a binary indicator for whether a hospital recorded any hospitalization due to gender-based violence for women older than 18. The sample consists of public hospitals located in Irish potato-suitable districts. The analysis examines the effects of the Irish potato harvest (the event of interest) while also suggesting no systematic changes occur during the coffee harvest season in potato-suitable districts. The Rwanda Seasonal Agricultural Surveys define Season C (July–September) as one of Rwanda's agricultural seasons, with Irish potatoes as one of its main crops and September as the harvest month (NISR 2016, NISR 2021). Calendar months are used, with February as the reference month. The regression controls for hospital fixed effects, month fixed effects, and province-by-harvest year fixed effects. Standard errors are clustered at the district level. *** p<0.01, ** p<0.05, * p<0.1.

Table A30: Effect of the Irish Potato Harvest on Monthly Food Consumption Expenditure in Potato-Suitable Districts

	(1) Monthly Food Consumption Expenditure
Month 1	0.03
	(0.03)
Month 2	0.00
	(.)
Month 3	-0.04**
	(0.02)
Month 4	-0.01
	(0.03)
Month 5	-0.04*
	(0.02)
Month 6	-0.01
	(0.02)
Month 7	-0.05**
	(0.02)
Month 8	-0.03
	(0.04)
Month 9	0.06***
	(0.01)
Observations	18457
Dependent variable mean	7.25

Notes: This table reports coefficients from the regression of monthly household food consumption expenditure on calendar-month dummies in potato-suitable districts. The dependent variable, drawn from the Integrated Household Living Conditions Survey (EICV), is measured in logs. The analysis examines the effects of the Irish potato harvest (the event of interest), while also suggesting no systematic changes occur during the coffee harvest season in potato-suitable districts. The sample consists of ever-married women (whose partner or most recent partner is alive) with an agricultural occupation residing in these districts. As in the previous consumption analysis, I restricted the sample to women with an agricultural occupation, as the variation in the data is tied to a harvest season. The event of interest is the Irish potato harvest, which takes place in September. The Rwanda Seasonal Agricultural Surveys define Season C (July–September) as one of Rwanda's agricultural seasons, with Irish potatoes as one of its main crops and September as the harvest month (NISR 2016, NISR 2021). Calendar months are used, with February as the reference month. The regression controls for women's education, age-group fixed effects (measured in five-year age groups), month fixed effects, district fixed effects, and harvest year fixed effects. Standard errors are clustered at the district level. *** p<0.01, ** p<0.05, * p<0.1.

A1.4 Additional Tables for Section 7: Robustness Checks

Table A31: Dynamic Impact of a Mill Opening on Ever-Married Women's Current Marital Status (Divorced or Separated) Among Pre-Mill Marriages

	Current Marital Status: Divorced or Separated	
	(1) First Marriage Began Before Mill Opening (Includes Women with Multiple Unions)	(2) First Marriage Began Before Mill Opening (One Lifetime Union)
Mill x Event Time -7- -2	0.00	0.01
	(0.02)	(0.02)
Mill x Event Time -1	0.00	0.00
	(.)	(.)
Mill x Event Time 0-1	0.01	-0.01
	(0.02)	(0.02)
Mill x Event Time 2-4	0.01	0.02
	(0.01)	(0.01)
Mill x Event Time 5-7	-0.01	0.00
	(0.02)	(0.02)
Observations	18276	16261
Dependent variable mean	0.13	0.12

Notes: This table reports coefficients from regressions of women's current marital status being divorced or separated on the interaction between residing in a mill's catchment area and event-year dummies, using equation (1). Column 1 presents results for ever-married women whose first marriage took place before the mill opening and includes women with multiple unions. Column 2 presents results for ever-married women whose first marriage took place before the mill opening and restricts the sample to women who married only once in their lives. Because the year of marriage is observed only for the first union, Column 2 provides the closest approximation in the data to a sample of women who married before the mill opening and remained married at the time of treatment. The dependent variable, drawn from the Demographic Health Surveys (DHS), is a binary indicator for whether a woman is currently divorced or separated at the time of the survey. The sample consists of ever-married women whose partner or most recent partner is alive. Event time is measured in years relative to the mill opening, with 0 indicating the year of the first mill's opening in the catchment area, negative values representing years before, and positive values representing years after. The omitted category is one year before the mill opening. Event time years are grouped to ensure a balanced number of observations across years. All regressions control for individual characteristics, including women's religion, education, agricultural occupation, age-group fixed effects (five-year age groups), survey month and year fixed effects, year of first marriage fixed effects, sub-district fixed effects, and district-by-harvest-year fixed effects. The catchment area is defined using a 4 km buffer around the mills. Standard errors are clustered at the sub-district level. I used the minimum and maximum event years –7 and 7 when using the divorce outcome variable, as in the household decision-making analysis, given that I discuss these two variables in parallel in Section 6.1. I used these specific minimum and maximum event years (instead of –6 and 6, as in the main results) because the household decision-making variables are defined for married women only, resulting in a smaller sample size compared to the sample in the main results. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A32: Dynamic Impact of a Mill Opening on Ever-Married Women Working and Working for Cash in the Past 12 Months (Excluding Year of First Marriage Fixed Effects)

	Work-Related Variables Measured Over the Past 12 Months	
	(1) Working	(2) Working for Cash
Mill x Event Time -6- -3	0.03	-0.01
	(0.02)	(0.03)
Mill x Event Time -2	0.02	0.01
	(0.02)	(0.04)
Mill x Event Time -1	0.00	0.00
	(.)	(.)
Mill x Event Time 0-1	0.02	-0.03
	(0.02)	(0.03)
Mill x Event Time 2-3	0.01	0.06**
	(0.02)	(0.03)
Mill x Event Time 4-6	0.01	0.06**
	(0.02)	(0.03)
Observations	21052	18352
Dependent variable mean	0.88	0.37

Notes: This table reports coefficients from regressions of work-related outcomes on the interaction between residing in a mill's catchment area and event-year dummies, using equation (1), except that year of first marriage fixed effects are excluded. The dependent variables, drawn from the Demographic Health Surveys (DHS), measure whether a woman worked in the past 12 months (Column 1) and whether she worked for cash in the past 12 months (Column 2). The sample for working for cash is restricted to women who reported working in the past 12 months, as defined by the survey question, which explains the smaller number of observations in Column 2. The sample consists of ever-married women whose partner or most recent partner is alive. The event of interest is the opening of a mill. Event time is measured in years relative to the mill opening, with 0 indicating the year of the first mill's opening in the catchment area, negative values representing years before, and positive values representing years after. The omitted category is one year before the mill's opening. Event time years are grouped to ensure a balanced number of observations across years. All regressions control for individual characteristics, including women's religion, education, age-group fixed effects (measured in five-year age groups), survey month and year fixed effects, and sub-district fixed effects, and district-by-harvest year fixed effects. In Column 2, I control for agricultural occupation (Since I later show that agricultural occupation does not change with treatment, it is considered a predetermined characteristic/variable and can therefore be used as a control.). I cannot control for agricultural occupation when 'worked in the past 12 months' is the dependent variable, given that the occupation information is not collected for women who did not work in the past 12 months. This is the only dependent variable for which I cannot control for agricultural occupation in my DHS analysis. The catchment area is defined using a 4 km buffer around the mills. Standard errors are clustered at the sub-district level. *** p<0.01, ** p<0.05, * p<0.1.

Table A33: Dynamic Impact of a Mill Opening on Ever-Married Women Experiencing Self-Reported Domestic Violence in the Past 12 Months: Physical or Sexual, Physical Only, Sexual Only (Excluding Year of First Marriage Fixed Effects)

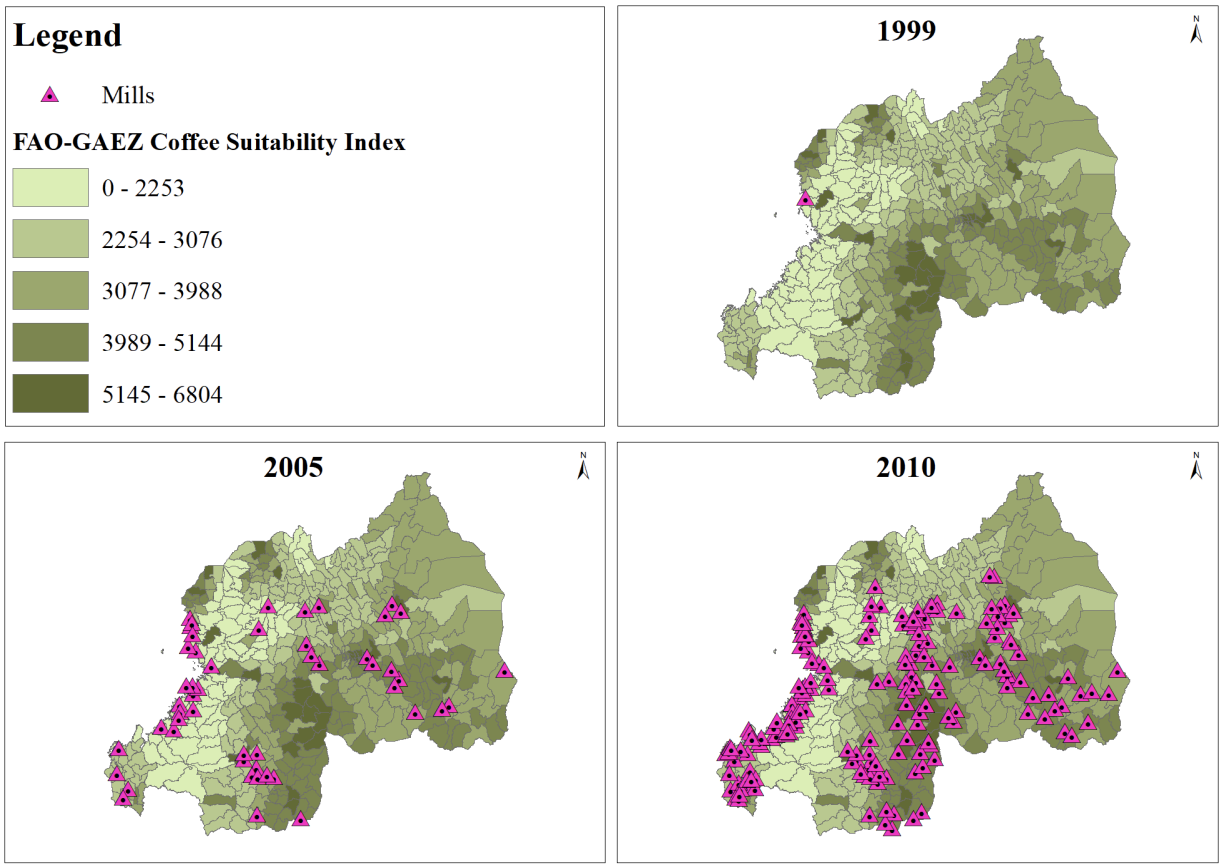
	Experiencing Domestic Violence in the Past 12 Months		
	(1) Physical or Sexual	(2) Physical	(3) Sexual
Mill x Event Time -6- -2	-0.02	0.01	-0.02
	(0.04)	(0.03)	(0.03)
Mill x Event Time -1	0.00	0.00	0.00
	(.)	(.)	(.)
Mill x Event Time 0-1	-0.12***	-0.08**	-0.07**
	(0.04)	(0.04)	(0.03)
Mill x Event Time 2-3	-0.10**	-0.09**	-0.05**
	(0.04)	(0.04)	(0.03)
Mill x Event Time 4-6	-0.07	-0.08*	-0.03
	(0.05)	(0.04)	(0.03)
Observations	5992	5991	5973
Dependent variable mean	0.32	0.29	0.11

Notes: This table reports coefficients from regressions of self-reported domestic violence outcomes on the interaction between residing in a mill's catchment area and event-year dummies, using equation (1), except that year of first marriage fixed effects are excluded. The dependent variables, drawn from the Demographic Health Surveys (DHS), measure whether a woman experienced physical or sexual violence (Column 1), physical violence only (Column 2), and sexual violence only (Column 3) in the past 12 months. The domestic violence module was administered to a randomly selected subsample of the full DHS sample, which explains the smaller number of observations compared to the work variables. Within this subsample, the analysis includes ever-married women whose partner or most recent partner is alive. Partnered women answer the questions for their current partners, while divorced or separated women answer them for their most recent partners. The event of interest is the opening of a mill. Event time is measured in years relative to the mill opening, with 0 indicating the year of the first mill's opening in the catchment area, negative values representing years before, and positive values representing years after. The omitted category is one year before the mill's opening. Event time years are grouped to ensure a balanced number of observations across years. All regressions control for individual characteristics, including women's religion, education, agricultural occupation, age-group fixed effects (measured in five-year age groups), survey month and year fixed effects, year of first marriage fixed effects, sub-district fixed effects, and district-by-harvest year fixed effects. The catchment area is defined using a 4 km buffer around the mills. Standard errors are clustered at the sub-district level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

A2 Additional Figures

A2.1 Additional Figures for Section 2: Institutional Context

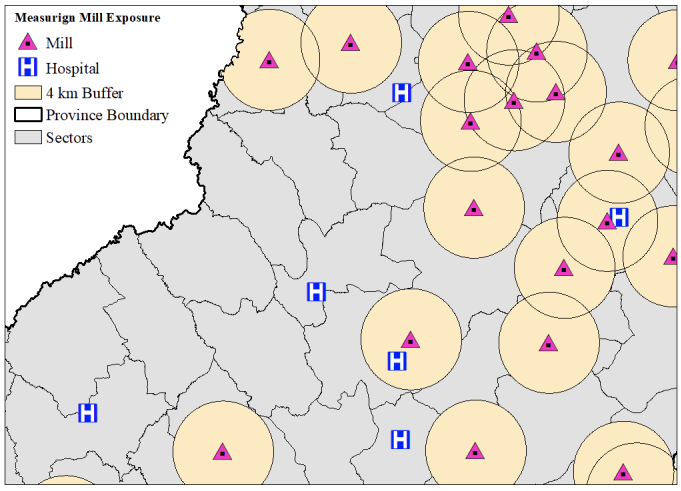
Figure A1: FAO-GAEZ Coffee Suitability Index and Expansion of Mills in Rwanda



Source: Harmonized data from Rwandan GeoPortal, Macchiavello and Morjaria (2020), FAO-GAEZ dataset (coffee suitability scores) and the Rwandan GeoPortal.
Notes: This figure visualizes the expansion of coffee mills in Rwanda over time in relation to the FAO-GAEZ Coffee Suitability Index. The maps display the geographic distribution of mills (represented by pink triangles) over three time points: 1999 (top-right), 2005 (bottom-left), and 2010 (bottom-right). The underlying shading represents coffee suitability scores, with darker shades indicating higher suitability.

A2.2 Additional Figures for Section 3: Data

Figure A2: Visualization of the Treated and Control Hospitals

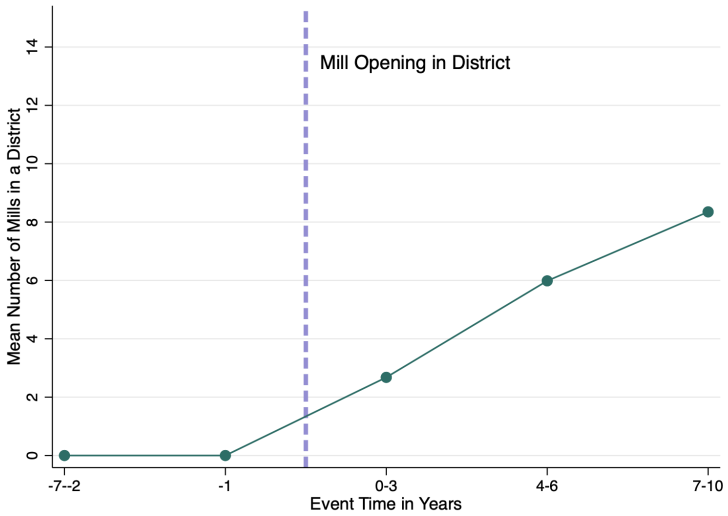


Source: Harmonized data from the Rwanda GeoPortal, Macchiavello and Morjaria (2020) and Hospital Management Information System (HMIS).

Notes: This map illustrates the treated and control hospitals used in the analysis. Treated hospitals are located within the 4 km radius catchment areas around mills, represented by the orange circular buffers, while control hospitals are outside the catchment areas but within the same province. Control hospitals are defined in this way due to the limited number of hospitals in Rwanda and their distribution. The analysis assumes that households within a mill's catchment area are more likely to visit treated hospitals, while households outside the catchment area are more likely to visit control hospitals due to proximity.

A2.3 Additional Figures for Section 4: Empirical Specification

Figure A3: Mean Number of Mills in a District by Event-Time (First Mill Opening in the District as the Event)

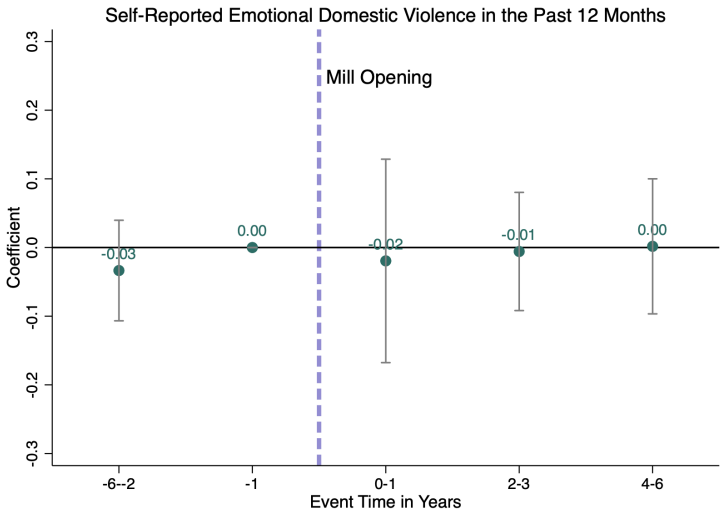


Source: Harmonized data from the Rwanda GeoPortal, Macchiavello and Morjaria (2020) and Integrated Household Living Conditions Survey (EICV)

Notes: This figure plots the mean number of mills in a district by event time, considering only districts with at least one mill by the end of the EICV sample period, 2017. The event of interest is the opening of a mill, marked by the purple dashed line. Event time is measured in years, with 0 indicating the year of the first mill's opening in the district, negative values representing years before, and positive values representing years after. The data is constructed using the EICV survey, which is not geocoded. First, a district is a larger geographic unit than a catchment area or a sub-district. Second, this figure captures the increasing number of mills within a district over time, which elevates the likelihood that individuals in the district are exposed to a mill. Based on these two points, the mill's effects are expected to become evident a few years post-opening, rather than immediately in year zero.

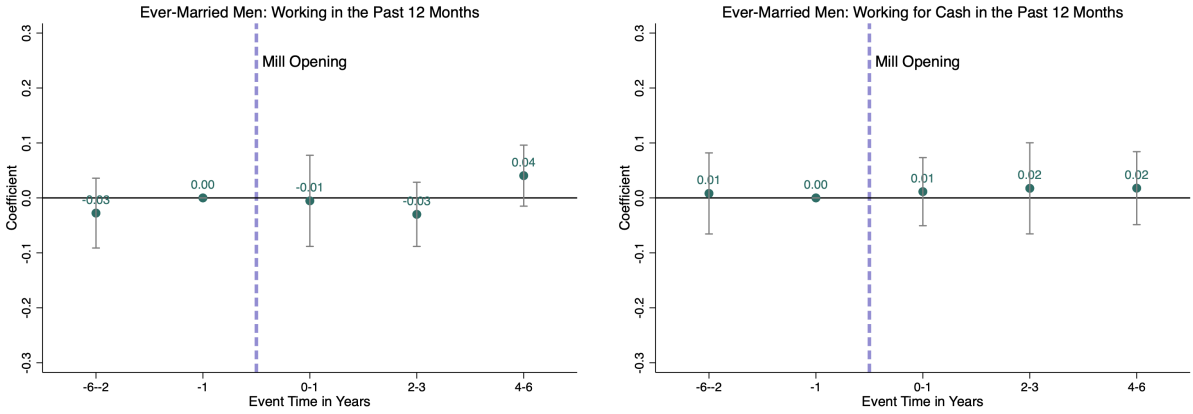
A2.4 Additional Figures for Section 5: Results

Figure A4: Dynamic Impact of a Mill Opening on Ever-Married Women's Self-Reported Domestic Violence in the Past 12 Months: Emotional Only



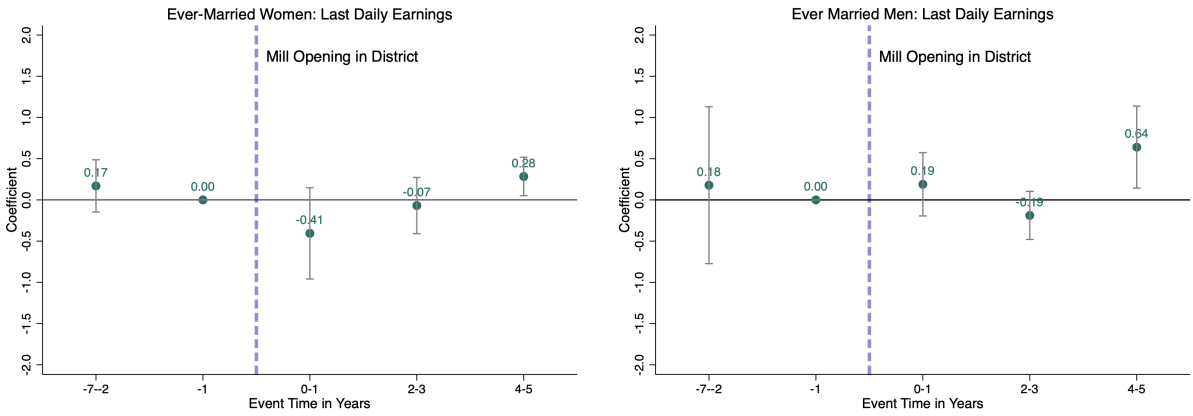
Notes: This figure plots the coefficients from the regression of self-reported emotional domestic violence on the interaction between residing in a mill's catchment area and event-year dummies, using equation (1). The dependent variable is drawn from the Demographic Health Surveys (DHS). The event of interest is the opening of a mill, marked by the purple dashed line. Event time is measured in years, where 0 indicates the year of the mill's opening, negative values represent years before, and positive values represent years after. Event time years are grouped to ensure a balanced number of observations across years. The domestic violence module was administered to a randomly selected subsample of the full DHS sample. Within this subsample, I use a sample of ever-married women whose partner or most recent partner is alive. Partnered women answer the questions for their current partners, while divorced or separated women answer them for their most recent partners. Unlike the main results and mechanisms, this analysis excludes the 2010 DHS cycle, as emotional domestic violence was not collected in that year. The y-axis shows the change in the probability of the outcome relative to one year before the mill's opening (the omitted category). Standard errors are clustered at the sub-district level, and 95% confidence intervals are displayed.

Figure A5: Dynamic Impact of a Mill Opening on Ever-Married Men's Working and Working for Cash in the Past 12 Months



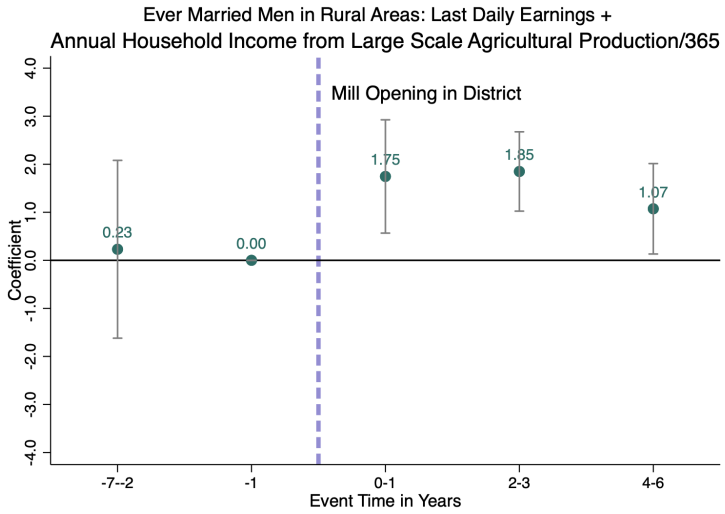
Notes: This figure plots the coefficients from the regression of men's working (left) and working for cash (right) on the interaction between residing in a mill's catchment area and event-year dummies, using equation (1). The dependent variables are drawn from the Demographic Health Surveys (DHS). The event of interest is the opening of a mill, marked by the purple dashed line. Event time is measured in years, where 0 indicates the year of the mill's opening, negative values represent years before, and positive values represent years after. Event time years are grouped to ensure a balanced number of observations across years. The sample consists of ever-married men whose partner or most recent partner is alive. The majority of these men are the husbands of currently married women in the DHS sample, aligning closely with the female sample structure. The y-axis shows the change in the probability of the outcome relative to one year before the mill's opening (the omitted category). Standard errors are clustered at the sub-district level, and 95% confidence intervals are displayed.

Figure A6: Dynamic Impact of a Mill Opening on the Last Daily Earnings of Ever-Married Women and Men



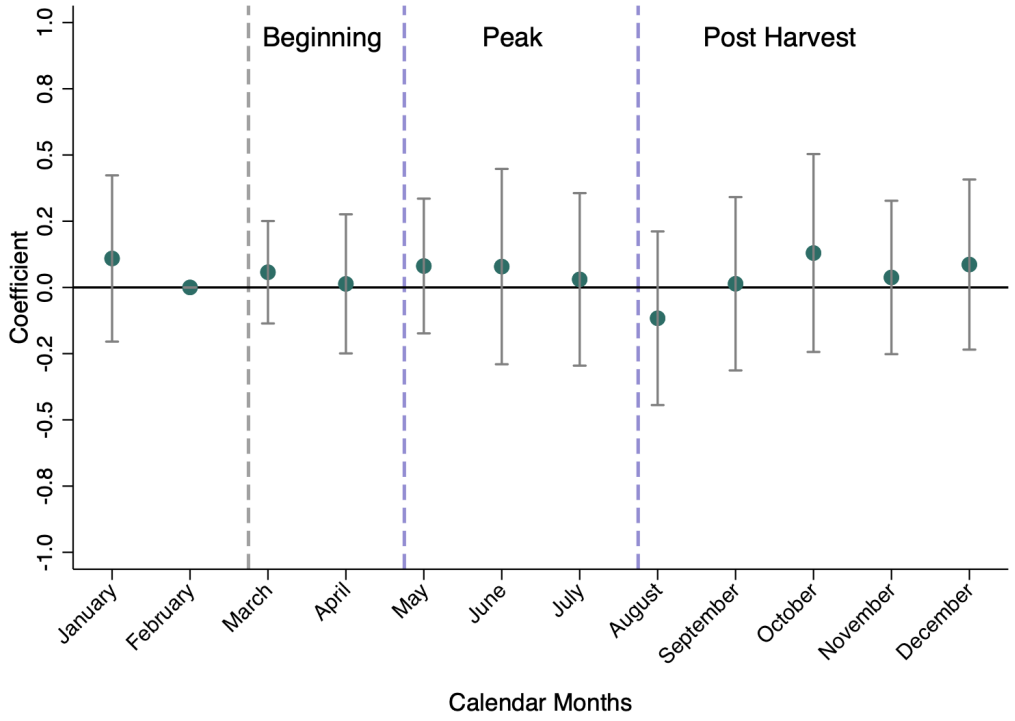
Notes: This figure plots the coefficients from the regression of women's and men's last daily earnings on the interaction between residing in a district with a mill and event-year dummies, using equation (2).The dependent variables, drawn from the Integrated Household Living Conditions Survey (EICV), are estimated using a Poisson pseudo-likelihood regression. The event of interest is the opening of a mill, marked by the purple dashed line. Event time is measured in years, with 0 indicating the year of the mill's opening, negative values representing years before, and positive values representing years after. Event time years are grouped to ensure a balanced number of observations across years. The women's sample consists of ever-married women whose partner or most recent partner is alive. The men's sample consists primarily of the husbands of married women in the sample, along with divorced or separated men whose most recent partner is alive. The y-axis shows the change in earnings relative to one year before the mill's opening (the omitted category). Standard errors are clustered at the district level, and 95% confidence intervals are displayed. The sample excludes individuals residing in Kigali, the capital, due to its highly urbanized labor market, where earnings levels differ significantly from the rest of the country. Including Kigali yields similar results, though confidence intervals for insignificant estimates in some event years are larger.

Figure A7: Dynamic Impact of a Mill Opening on Earnings of Ever-Married Men in Rural Areas



Notes: This figure plots the coefficients from the regression of men's earnings in rural areas (last daily earnings + their annual household income from large-scale agricultural production as defined in the dataset divided by 365) on the interaction between residing in a district with a mill and event-year dummies, using equation (2). The dependent variable, drawn from the Integrated Household Living Conditions Survey (EICV), is estimated using a Poisson pseudo-likelihood regression. The event of interest is the opening of a mill, marked by the purple dashed line. Event time is measured in years, with 0 indicating the year of the mill's opening, negative values representing years before, and positive values representing years after. Event time years are grouped to ensure a balanced number of observations across years. The sample consists of ever-married men whose partner or most recent partner is alive and reside in rural areas. The y-axis shows the change in earnings relative to one year before the mill's opening (the omitted category). Standard errors are clustered at the district level, and 95% confidence intervals are displayed.

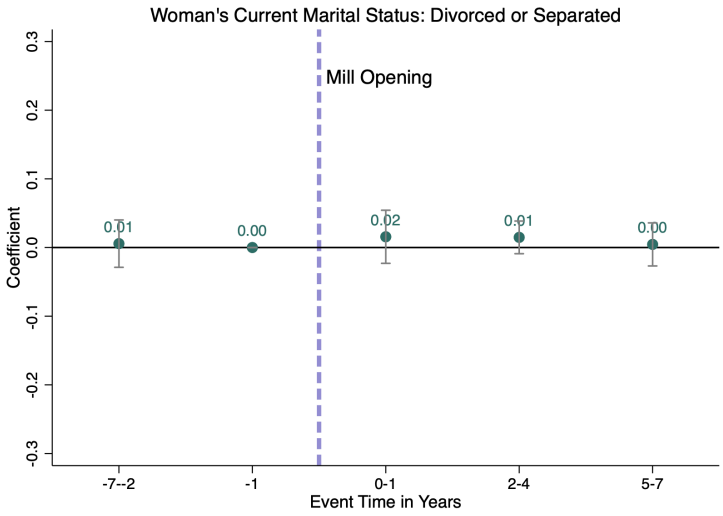
Figure A8: Dynamic Impact of the Beginning of Mill Operations on Hospitalizations for Bone and Joint Disorders other than Fractures for Women Older than 18 (Non-Gender Based Violence Hospitalizations)



Notes: This figure plots the coefficients from the regression of hospitalizations for bone and joint disorders other than fractures (e.g., osteoarthritis, gout, rheumatoid arthritis, lupus, bursitis) for women older than 18 on the interaction between a hospital being in a mill's catchment area and calendar-month dummies, using equation (3). The dependent variable is drawn from the Hospital Management Information System (HMIS) (2012-2015), and is not recorded as gender-based violence in the hospital records. The variable is measured in logs. The event of interest is the start of mill operations, which coincides with the beginning of the coffee harvest season, marked by the gray dashed line (March). The harvest season peaks between May and July and is followed by the post-harvest period from August to December, when mill operations have ended. Calendar months are used, with February as the reference month. The y-axis shows the change in the probability of hospitalization relative to February, the month before mill operations begin. Standard errors are clustered at the district level, and 95% confidence intervals are displayed.

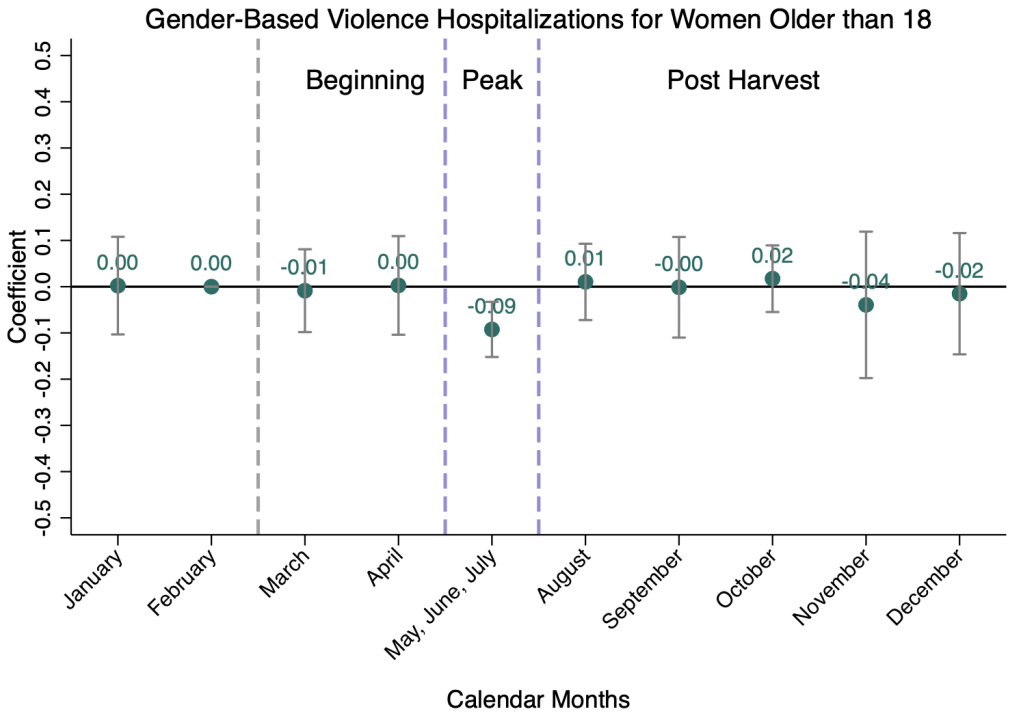
A2.5 Additional Figures for Section 7: Mechanisms

Figure A9: Dynamic Impact of a Mill Opening on Ever-Married Women's Current Marital Status Being Divorced or Separated



Notes: This figure plots the coefficients from the regression of women's current marital status being divorced or separated on the interaction between residing in a mill's catchment area and event-year dummies, using equation (1). The dependent variables are drawn from the Demographic Health Surveys (DHS). The event of interest is the opening of a mill, marked by the purple dashed line. Event time is measured in years, where 0 indicates the year of the mill's opening, negative values represent years before, and positive values represent years after. Event time years are grouped to ensure a balanced number of observations across years. The sample consists of ever-married women whose partner or most recent partner is alive. The y-axis shows the change in the probability of the outcome relative to one year before the mill's opening (the omitted category). Standard errors are clustered at the sub-district level, and 95% confidence intervals are displayed.

Figure A10: Dynamic Impact of the Beginning of Mill Operations on Physical or Sexual Gender-Based Violence Hospitalizations for Women Older than 18 (Proxy for Domestic Violence): Peak Harvest Months Pooled

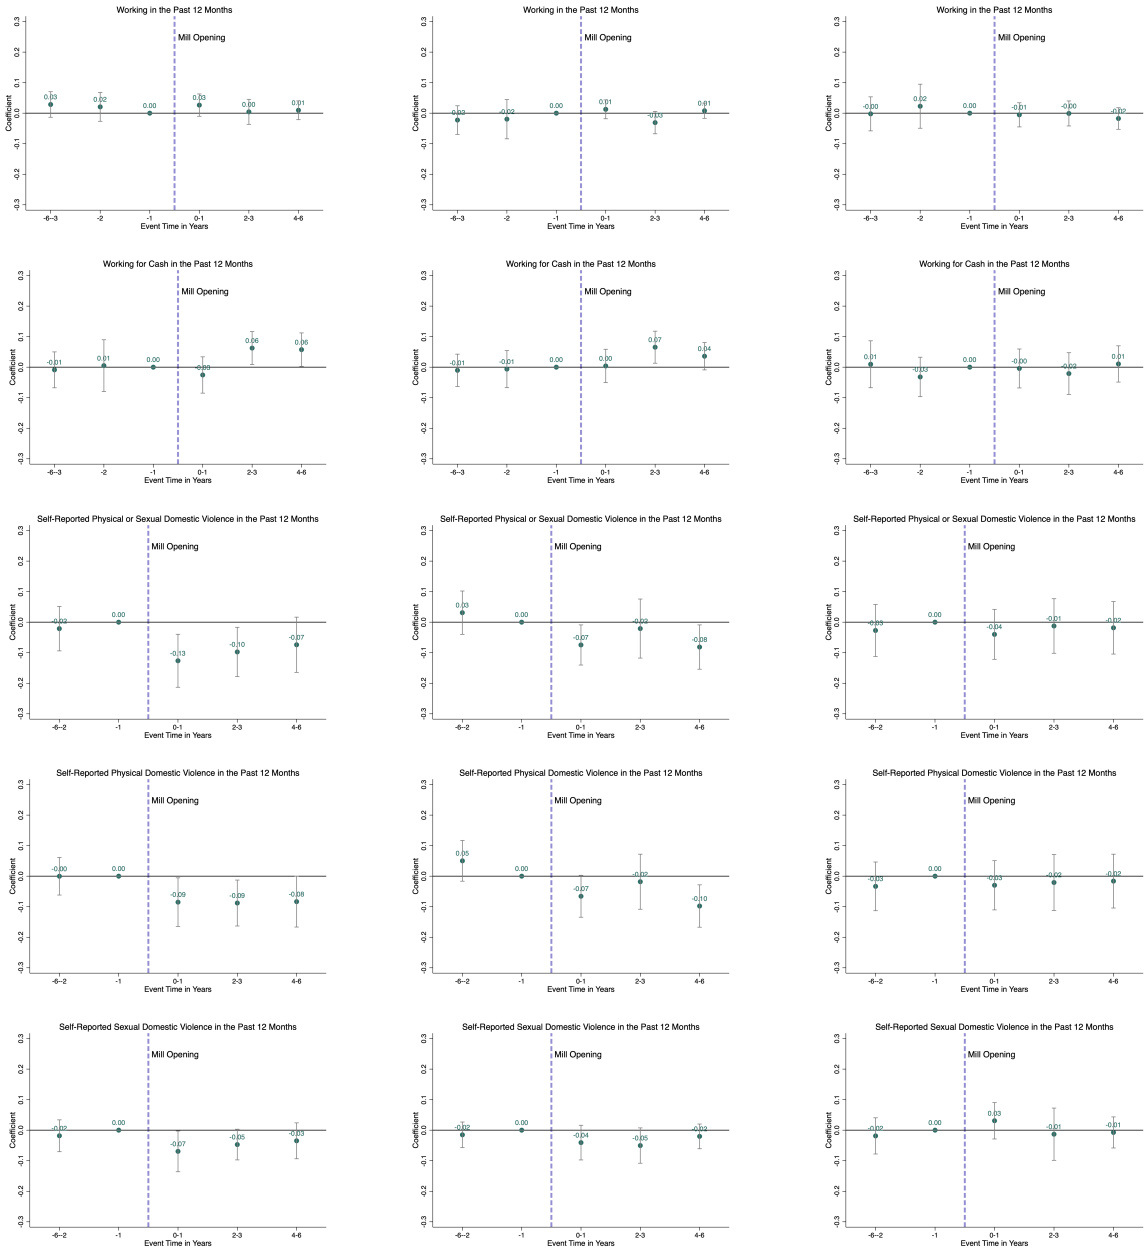


Notes: This figure plots the coefficients from the regression of hospitalizations for physical or sexual gender-based violence for women older than 18 on the interaction between a hospital being in a mill's catchment area and calendar-month dummies, using equation (3). The dependent variable, drawn from the Hospital Management Information System (HMIS) as a proxy for domestic violence, is a binary indicator reflecting whether a hospital recorded any hospitalization due to gender-based violence among women older than 18. The event of interest is the start of the mill's operations, which coincides with the beginning of the coffee harvest season, marked by the gray dashed line (March). The harvest season peaks between May and July, and is followed by the post-harvest period from August to December, when mill operations have ended. In this specification, the peak months of the harvest (May, June, and July) are pooled together to assess the overall impact of the period when labor demands and the cost of physically incapacitating one's wife are at their highest. Calendar months are used, with February as the reference month. The sample consists of 43 public hospitals in Rwanda that recorded monthly gender-based violence cases during 2012-2019, out of the 48 public hospitals in the country by 2019, representing the universe of such hospital records in the country. The y-axis shows the change in the probability of hospitalization relative to February, the month before mill operations begin. Standard errors are clustered at the district level, and 95% confidence intervals are displayed.

A2.6 Additional Figures for Section 8: Robustness Checks

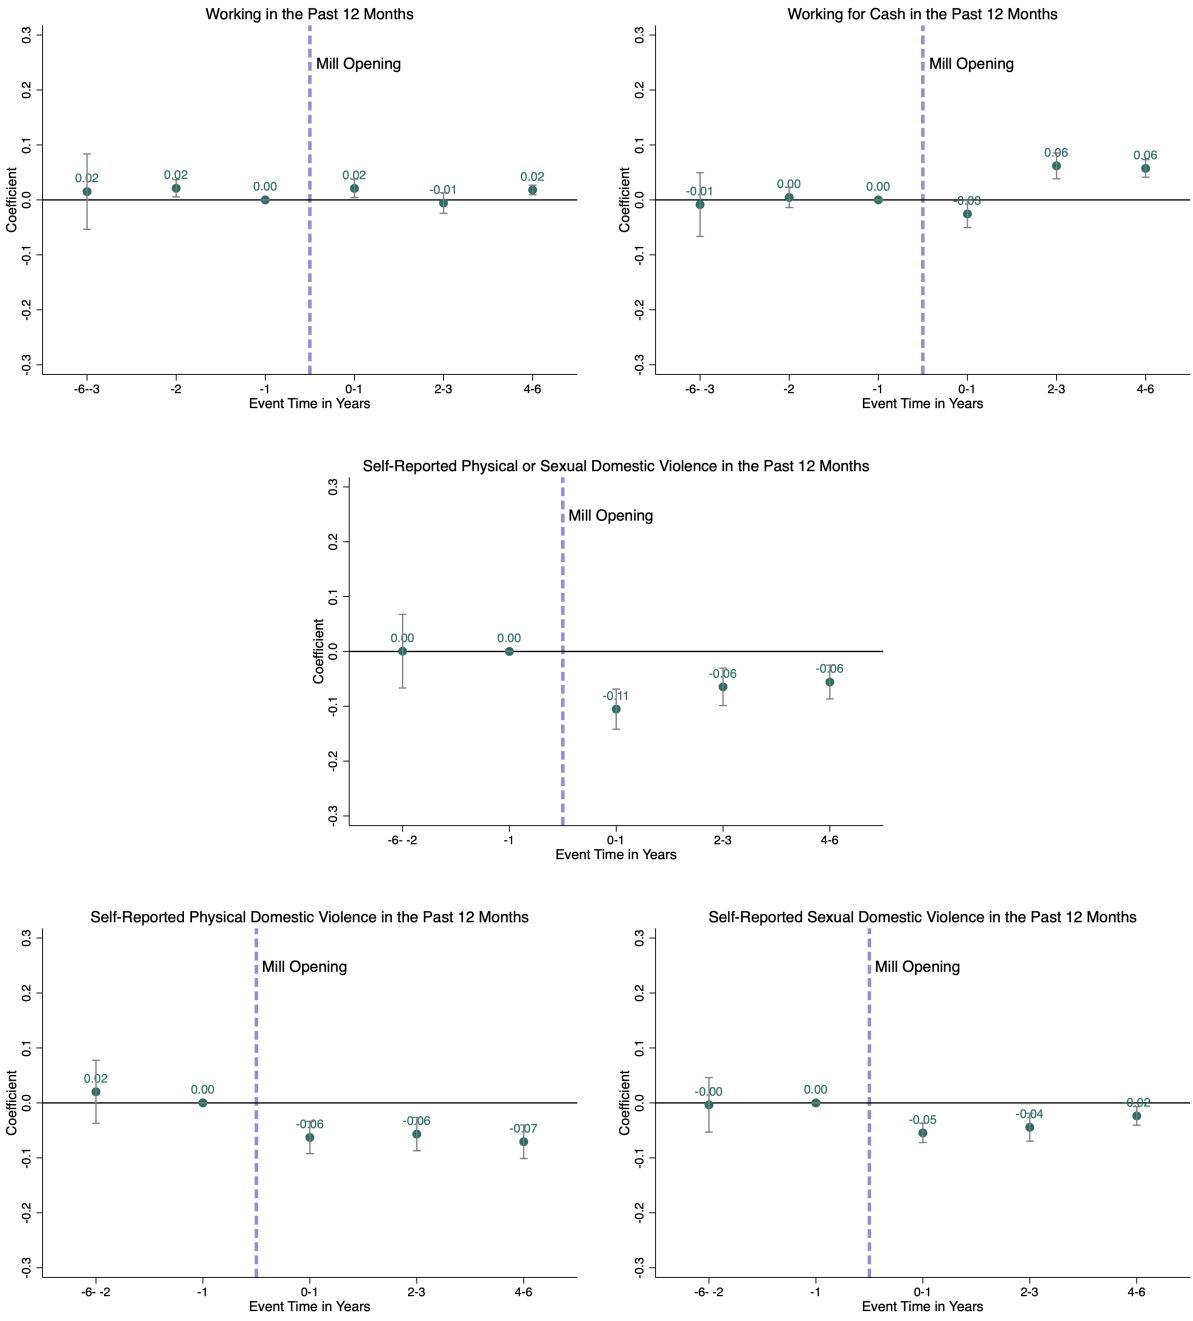
Figure A11: Effect of a Mill Opening on Ever-Married Women's Working, Working for Cash and Self-Reported Domestic Violence in the Past 12 Months Using Different Buffer Radii

Column 1: Buffer Radius is 4 km Column 2: Buffer Radius is 5 km Column 3: Buffer Radius is 10 km



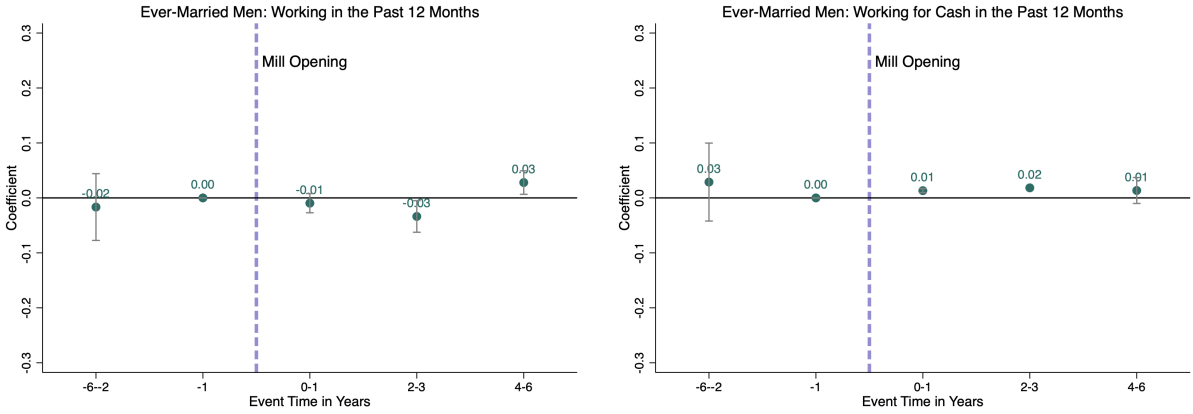
Notes: This figure presents robustness checks for the effect of a mill opening on women's working, working for cash, and self-reported domestic violence in the past 12 months. Column 1 replicates the main results, using a 4 km buffer radius to define a mill's catchment area. Columns 2 and 3 expand the buffer radius to 5 km and 10 km, respectively, to test the sensitivity of the estimates to different definitions of exposure. The dependent variables are drawn from the Demographic Health Surveys (DHS), and regressions follow equation (1). The event of interest is the opening of a mill, marked by the purple dashed line. Event time is measured in years, where 0 indicates the year of the mill's opening, negative values represent years before, and positive values represent years after. Event time years are grouped to ensure a balanced number of observations across years. The y-axis shows the change in the probability of the outcome relative to one year before the mill's opening (the omitted category). Standard errors are clustered at the sub-district level, and 95% confidence intervals are displayed.

Figure A12: Dynamic Impact of a Mill Opening on Ever-Married Women's Working, Working for Cash and Experiencing Self-Reported Domestic Violence in the Past 12 Months: Physical or Sexual, Physical Only, Sexual Only in the Past 12 Months using Sun and Abraham (2020)



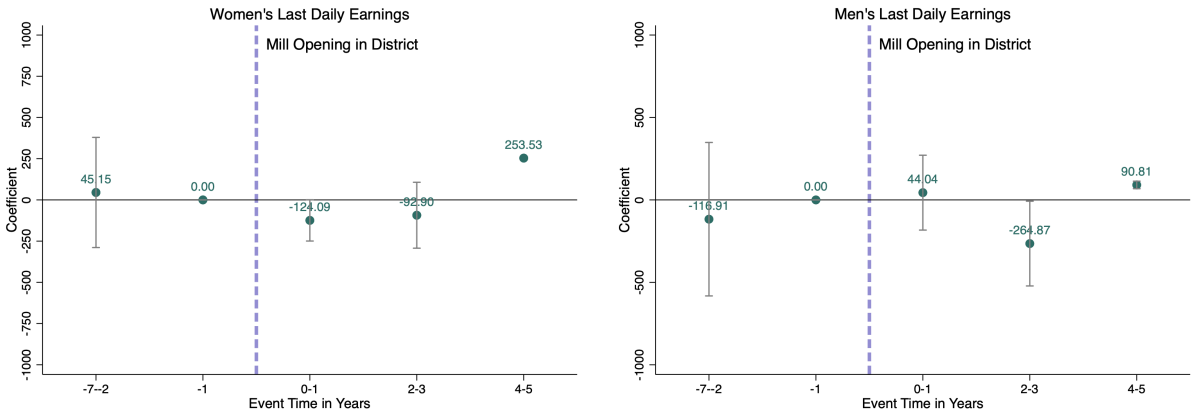
Notes: This figure replicates the main analysis using the method proposed by Sun and Abraham (2020). The sample, outcome variables, and event-year structure remain the same as in Figures 4 and 5. Standard errors are clustered at the sub-district level, and 95% confidence intervals are displayed.

Figure A13: Dynamic Impact of a Mill Opening on Ever-Married Men's Working and Working for Cash in the Past 12 Months using Sun and Abraham (2020)



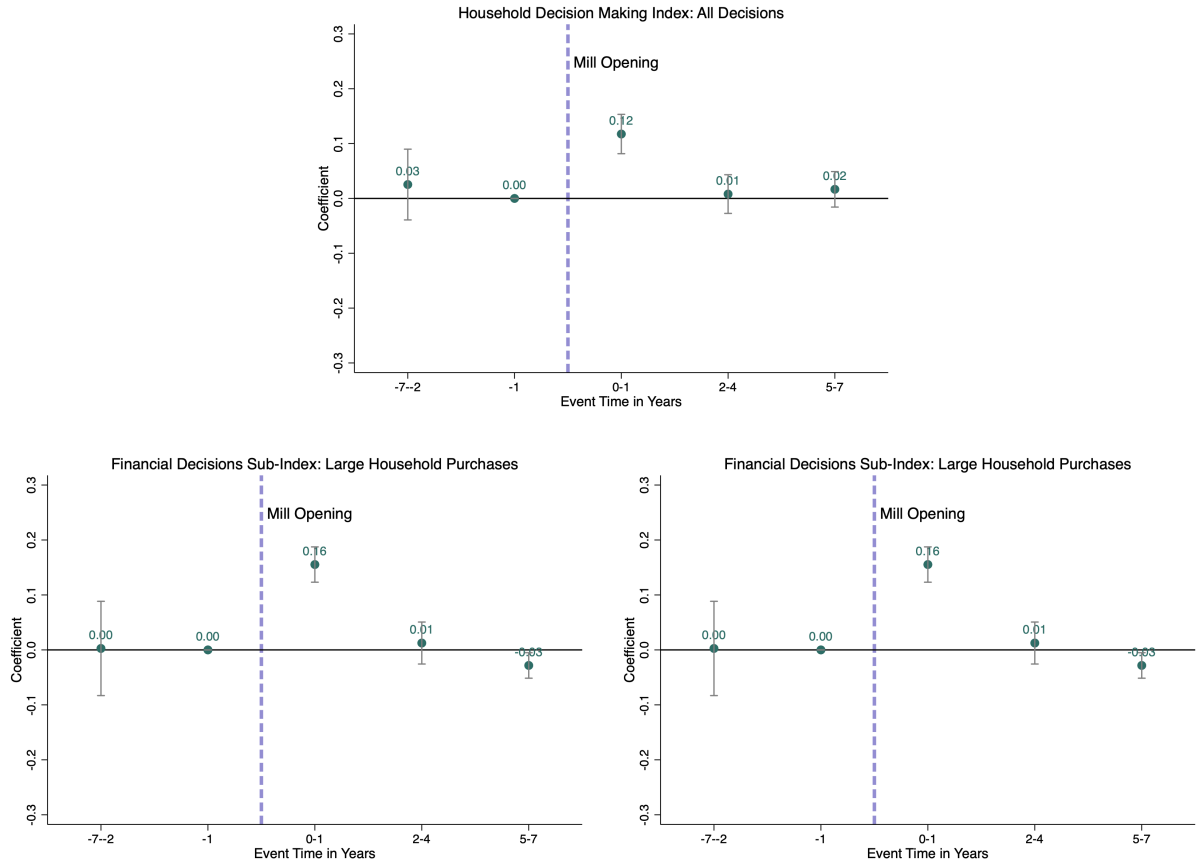
Notes: This figure replicates the main analysis using the method proposed by Sun and Abraham (2020). The sample, outcome variables, and event-year structure remain the same as in Figure A5. Standard errors are clustered at the sub-district level, and 95% confidence intervals are displayed.

Figure A14: Dynamic Impact of a Mill Opening on Ever-Married Women's and Men's Last Daily Earnings using Sun and Abraham (2020)



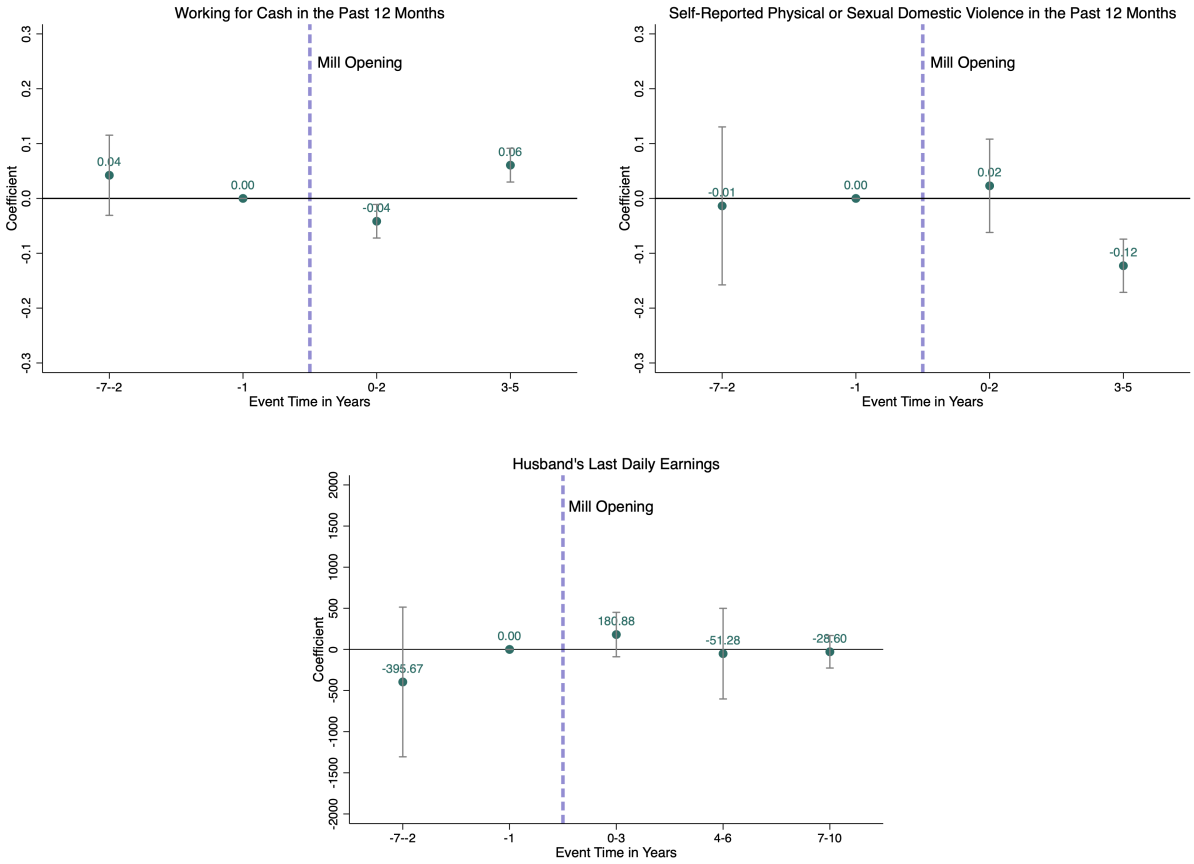
Notes: This figure replicates the main analysis using the method proposed by Sun and Abraham (2020). The sample, outcome variables, and event-year structure remain the same as in Figure A6. Unlike the main analysis, this regression does not use a Poisson pseudo-likelihood approach, as the code for the estimator from Sun and Abraham (2020) is not built for it. Instead, the regression is estimated using a linear model. As in the main analysis, individuals not working remain coded as 0. Standard errors are clustered at the district level, and 95% confidence intervals are displayed.

Figure A15: Dynamic Impact of a Mill Opening on Married Women's Participation in Household Decision Making (Alone or Jointly with the Husband): All, Financial and Non-Financial Decisions using Sun and Abraham (2020)



Notes: This figure replicates the mechanisms analysis using the method proposed by Sun and Abraham (2020). The sample, outcome variables, and event-year structure remain the same as in Figure 9. Standard errors are clustered at the sub-district level, and 95% confidence intervals are displayed.

Figure A16: Dynamic Impact of a Mill Opening on Ever-Married Women's Working for Cash and Self-Reported Domestic Violence in the Past 12 Months: Women with Different Occupations from Their Current or Most Recent Husband using Sun and Abraham (2020)



Notes: This figure replicates the mechanisms analysis using the method proposed by Sun and Abraham (2020). The sample, outcome variables, and event-year structure remain the same as in Figure 10. Standard errors are clustered at the sub-district level, and 95% confidence intervals are displayed.

A3 Data Sources

A3.1 Panel of Mills

First, I complement the mill data with spatial data from the Rwanda GeoPortal. The portal provides maps of various geographical boundaries, including sectors, districts, and provinces. Rwanda has 416 sub-districts, 30 districts, and 5 provinces, with a mean area of approximately 58.4 km^2, 810 km^2, and 5,552 km^2, respectively.

According to the 2002 Rwandan Population and Housing Census, the average sub-district population is 1,959, with an average female population of 1,024 at the sub-district level. The average district population is 27,698, with an average female population of 14,499 at the district level.

The final panel of mills contains information on each mill's location at different levels (including exact GPS coordinates), the number of coffee trees in the sub-district where the mill is located across different years, and the FAO-GAEZ coffee suitability score. The data also include details on mill characteristics, such as ownership type.

A3.2 Individuals and Household Level Data

Rwandan Demographic Health Surveys. DHS asks whether women are paid "cash only" or "cash and in-kind." Field surveys indicate that mills provide clothing and other in-kind benefits (including coffee) to women working in the mills in addition to a daily wage. Thus, I construct my cash earnings variable based on women working for both cash and in-kind, rather than cash only. According to the data, occupations that pay cash-only wages to women are predominantly non-agricultural (e.g., managers in Kigali, sales, etc.). Upon a mill opening, I find statistically significant changes in the probability of working for cash when I define the variable as women working for both cash and in-kind. However, when defining the variable as cash-only employment, there is no significant change in the probability of working for cash upon a mill opening. This further confirms that "cash and in-kind" correctly captures the earnings structure of women working in the mills. I perform the same analysis for husbands and find no change in their probability of working for cash under either definition, confirming that the expansion did not affect men's earnings type.

The domestic violence module is prepared with respect to WHO guidelines "Putting Women First: Ethical and Safety Recommendations for Research on Domestic Violence against Women" World Health Organization, 2001. Physical domestic violence behaviors include when the husband pushed, shook, threw something at the wife, slapped, twisted her arm, pulled her hair, punched with his fist or an object that could hurt her, kicked her, dragged her, beat her up, tried to choke her, burned her on purpose, or threatened or attacked her with a knife, gun, or any other weapon.

Sexual domestic violence behaviors include when the husband physically forced her to have sexual intercourse or perform other sexual acts she did not want to. Emotional domestic violence includes when the husband says or does something to humiliate the wife in front of others, threatens to hurt or harm her, or insults her or makes her feel bad about herself. Based on the described behaviors, physical domestic violence has the greatest potential to physically incapacitate the wife.

DHS data allow me to observe outcome variables for 2004, 2005, 2010/2011, 2014/2015, and 2019/2020. The 2005 cycle was collected from February to July 2005, the 2010/11 cycle from September 2010 to March 2011, the 2014/15 cycle from November 2014 to April 2015, and the 2019/20 cycle from November 2019 to July 2020. Given that the harvest season runs from March to July and key variables capture individuals' experiences in the past 12 months, the 2005 data cycle contains information on both the 2004 and 2005 harvest seasons. The 2010/11 cycle provides information on the 2010 harvest season, the 2014/15 cycle on the 2014 harvest season, and the 2019/20 cycle on the 2019 harvest season. Thus, DHS data allow me to observe a total of five harvest years: 2004, 2005, 2010, 2014, and 2019.

A3.3 Administrative Hospital Level Gender-Based Violence Data

The sample consists of 43 public hospitals in Rwanda, including district, provincial, and national referral hospitals, that recorded gender-based violence cases during the study period. These 43 hospitals provided the gender-based violence data used in this analysis, out of the 48 public hospitals in the country by 2019. The remaining five hospitals, which do not record such violence cases, are teaching hospitals or specialized hospitals (e.g., orthopedics) (MOH, 2020). More than 90% of Rwandans are covered by the government's mandatory health insurance system, with the poorest citizens benefiting from public health services free of charge (Lu et al. 2012, Koch et al. 2022).

In addition to the 48 public hospitals, there are eight private hospitals in the country as of 2019. Based on the institutional structure of the gender-based violence services provided in the country (One Stop Centers which provide medical, police and legal help for gender based violence victims under one roof and are free of charge), the services are in the public hospitals, not private.[1]

Most of the hospitals in the dataset are district hospitals, serving as the first point of hospital-level health services in a district in Rwanda. The patient is referred to either a provincial hospital (one geographical unit above the district) or a national referral hospital based on the severity of the patient's health condition. National referral hospitals are mostly located in Kigali (the capital) or other urban areas, equipped with the latest technology, and provide care under the public national health insurance system when referred by a district hospital or health clinic. Thus, they do not serve as the first point of care. Teaching hospitals focus primarily on medical research.

[1] Sviatschi and Trako (2021) finds that similar centers in Peru reduced gender-based violence and increased human capital investments in children, leading to higher school enrollment, attendance, and test scores.

A4 Supplementary Empirical Analysis

A4.1 What predicts mill placement?

The key identifying assumption of the empirical strategy is that the average outcomes for the treatment and control groups would have parallel trends in the absence of treatment, a mill placement. A mill opening at a specific location in a given year is assumed to be uncorrelated with other determinants of changes in the outcome variables (women's labor market outcomes and domestic violence) over time (the treatment is not endogenous). The parallel trends assumption is violated if a mill placement is endogenous. This is because, in counterfactual, the areas who are exposed to the mill would have diverged anyway, regardless of the mill.

To provide suggestive evidence in favor of my identifying assumption, I estimate, at the sub-district level, the determinants of having the first mill during the period when mills expanded rapidly, between 2005-2010, and having a mill by the end of my sample, 2018. Firstly, I focus on the variables related to coffee cultivation, including the historical number of coffee trees and the FAO-GAEZ coffee suitability index. I also include the shares of women who are unpaid family workers, self-employed, and completed primary school at the sub-district level before the expansion of the mills and the percentage of men who completed primary school in the regression. This is to confirm that mill placement is uncorrelated with factors that can affect the evolution of female paid employment and domestic violence rates in Rwanda. I used the Rwanda Population and Housing Census 2002 to create the variables. I also included the total population and female population at the sub-district level.

I also include other variables related to promoting female empowerment, including the share of women in a consensual union, polygamous marriage, without assets and the number of daughters per woman. If these variables are uncorrelated with mill placement, then it provides supporting evidence that the mills are not placed to promote female empowerment; hence, mill placement is not endogenous. Moreover, I include genocide intensity since it is also a variable that is found to impact the probability of experiencing domestic violence in Rwanda (LaMattina, 2017). It is also possible that the government or NGOs are more likely to financially support the opening of a mill in genocide intense areas to promote female empowerment, which can make the opening of a mill endogenous. District fixed effects are included in all specifications.

I report the results in Table A34. The results support the evidence that variables related to coffee, the number of coffee trees in 1999 and FAO-GAEZ coffee suitability predicts mills placement. Specifically, sub-districts with more coffee trees in 1999 are more likely to have their first mill between 2005-2010. Those areas are also more likely to have a mill by 2018. Similarly, sub-districts with a higher FAO-GAEZ coffee suitability index are more likely to have their first mill

between 2005-2010. Based on the results, the historical number of coffee trees is more likely to be correlated with the rapid expansion of the mill compared to the FAO-GAEZ coffee suitability index. This is not surprising. Before the expansion, not every coffee-suitable area in Rwanda had coffee trees and it takes a minimum of 3-5 years for a coffee tree to produce coffee cherries. Thus, it is likely to observe the rapid expansion of the mills between 2005-2010 in areas with coffee trees before the expansion, like in 1999.

All variables other than the number of coffee trees in 1999 and the FAO-GAEZ coffee suitability index are statistically insignificant. Specifically, pre-expansion female unpaid employment and female, male primary education rates do not predict mill placement. None of the variables related to female empowerment is statistically significant. Genocide-intensity does not have a statistically significant impact on mill expansion. I also test the joint significance of the relationship between a mill opening and non-coffee variables. For both regressions, I cannot reject the null hypothesis (p-value=0.4954 and p-value=0.6026). These results suggest that mills are not placed to promote female empowerment. They are uncorrelated to the factors that can impact the evolution of female paid employment and domestic violence over time. Thus, conditional on mill suitability based on coffee-related variables, the opening of a mill is assumed to be random.

Table A34: Sub-district Level Baseline Characteristics that Predict Mill Opening

	(1) First Mill in 2005-10	(2) Mill by 2018
Log Coffee Trees in 1999	0.03***	0.04***
	(0.01)	(0.01)
FAO-GAEZ Coffee Suitability Index	0.07*	0.03
	(0.04)	(0.04)
Log Population in 2002	-0.10	-0.14
	(0.49)	(0.49)
Log Female Population in 2002	0.07	0.20
	(0.49)	(0.49)
Share of Self-Employed Women in 2002	0.27	0.17
	(0.55)	(0.54)
Share of Unpaid Worker Women in 2002	0.00	0.44
	(0.58)	(0.58)
Share of Primary-Educated Women in 2002	1.45	1.38
	(1.09)	(1.08)
Share of Primary-Educated Men in 2002	-0.14	-0.40
	(1.20)	(1.19)
Number of daughters per Woman in 2002	-0.22	-0.13
	(0.22)	(0.22)
Share of Women in a Consensual Union in 2002	-1.20	-0.57
	(0.78)	(0.77)
Share of Women in a Polygamous Marriage in 2002	0.17	-1.80
	(1.94)	(1.94)
Share of Women without Assets in 2002	0.32	0.02
	(0.57)	(0.57)
Genocide Intensity Index at the Commune Level	-0.02	-0.04
	(0.04)	(0.03)
District FE	✓	✓
Number of Observations	348	348
Dependent variable mean	0.26	0.39
Adjusted R^2	0.24	0.39

Notes: This table presents linear regressions at the sub-district level, estimating the determinants of having the first mill during the rapid expansion period (2005–2010) in column (1) and having a mill by the end of the sample period (2018) in column (2). The dependent variable is an indicator for whether a mill was present in the sub-district during the respective time period. The FAO-GAEZ coffee suitability index and the genocide intensity index are both standardized. The data is at the sub-district level. District fixed effects are included. *** p<.01, ** p<.05, * p<.1.

A5 Theory Appendix

A5.1 Theoretical Framework

A mill opening increases the value of the wife's capacity to work, either via the husband selling the couple's harvest to the mill for a high premium and/or the wife working in the mill for pay which she shares her earnings with the husband. Thus, the conceptual framework studies how women's paid employment and their husband's earnings affect domestic violence.

A5.2 Setup

Preferences. The household consists of a wife and husband, $j \in \{w,h\}$. I assume that the preferences can be represented by the utility functions

$$U_h = yI_h + (1-x+d\tau)I_w + d\alpha(I) + \theta_h \quad and \quad U_w = (1-y)I_h + (x-d\tau)I_w + \theta_w - dv_w \tag{1}$$

where I_j for $j \in \{w,h\}$ indicates personal earnings for the wife and husband. I is total household resources where $I = I_w + I_h$. I_w is a function of the wife's productivity in the labor market, $I_w(\rho)$, where $\rho \in (0,1)$ is the productivity rate. I assume that the wife's earnings is increasing in her productivity. Incapacitation can be thought of ρ being very close to 0, which makes I_w very close to 0. I denote $I_w(\rho)$ as I_w for notational simplicity. In my context, the husband benefits from the wife's productivity in two ways. First, the husband's earnings is a function of the wife's productivity. Second, the wife shares a portion of her earnings with him. For simplicity, I incorporate only the second way into the framework below.[2]

According to Rwandan social norms, the husband controls household income and household decision-making (ICRW 2011, Bayisenge 2010, Ya-Bititi et al. 2019). The husband keeps yI_h of his personal earnings, $y \in (0,1)$, and gives $(1-y)I_h$ of it to the wife for her expenses. The wife keeps xI_w of her personal earnings, $x \in (0,1)$, and turns $(1-x)I_w$ of it to the husband. In my data, 65% of women decide how to spend their earnings jointly with their husbands, not by themselves, which supports the model. I assume that x and y are predetermined in the beginning of the marriage. Based on the setup, the productivity of the wife affects the husband's utility.

The main difference between the husband and the wife is that the husband can extract more money than his spouse turns over to him (based on the predetermined x value) by inflicting violence. Mainly, the husband can use violence to get a bigger share of her earnings and maximize his utility from marriage. Violence decision is parametrized by $d \in \{0,1\}$. $\tau \in (0,x]$ is the extrac-

[2] Again, for simplicity, I_h is not a function of his productivity since the wife's, not the husband's, incapacitation is the focus of the paper.

tion rate. If the husband chooses to inflict violence on his wife, $d = 1$. If he chooses not to be violent, $d = 0$. τ captures *instrumental* violence, i.e., the husband uses violence as a tool to extract resources from his wife.

The term $\alpha(I)$ denotes the husband's non-monetary utility from inflicting violence. I assume that $\alpha(I) > 0$. It captures *expressive* violence in the framework, which means that with the parameter, the benefit from domestic violence enters the husband's utility directly, not in an instrumental way (Gelles 1974, Tauchen et al. 1991). The parameter can be seen as the husband's taste for violence. It can be an intrinsic love for violence, or violence can represent a stress outlet. I assume that the husband's taste for violence is decreasing in household resources, which affects violence. Such an income effect could, for instance, stem from stress relief. In such a case, when there are more resources in the household, there is less financial stress and marginal utility from stress relief is lower. Thus violence decreases.

The term θ_j for $j \in \{w,h\}$ indicates the private level of satisfaction with the marriage for the wife and husband. As a key assumption of the model, each spouse's level of satisfaction with the marriage remains private information. θ_j follows distribution F_j with support $[\underline{\theta_j}, \overline{\theta_j}]$. $\underline{\theta_j}$ is high enough for the marriage to be intact. v_w is the disutility from violence for the wife where $v_w > 0$. U_w is decreasing in domestic violence and increasing in her and the husband's earnings and her private level of satisfaction with the marriage. U_h is increasing in his and his wife's earnings and his private level of satisfaction with the marriage.

I assume that their outside options can be represented by the utility functions

$$S_h = I_h \quad and \quad S_w = I_w. \tag{2}$$

Each spouse only enjoys her/his personal earnings as an outside option. I define the outside option as the utility of being single (divorced).[3] Due to paid employment, the wife has personal earnings to support herself financially if she initiates divorce.[4]

Timing. The husband keeps a portion of his personal income to himself and gives the remaining portion to the wife for her expenses. The wife keeps a portion of her personal income and turns the rest to her husband. The husband observes the resources he receives from his wife, $(1-x)I_w$, and either chooses to inflict violence to extract more resources from the wife, τI_w, or he chooses

[3]Divorce is a credible threat in the Rwandan context. In 2008, Rwanda became the first country in Sub-Saharan Africa to criminalize all forms of domestic violence and allow women to divorce their husbands unilaterally if their husbands are violent towards them. Sanin (2024) shows that women more likely to be in violent marriages are more likely to get divorced after the law.

[4]I conceptualize the utility of being single with personal earnings only as a simplification. A paid job also has non-monetary benefits that can constitute a part of the wife's outside option. For example, due to having a job, the wife has a social network that can provide her with a place to stay when she is separated.

not to be violent and remains with $(1-x)I_w$. If he chooses to inflict violence, his utility becomes $U_h = yI_h + (1-x+\tau)I_w + \alpha_h(I_h+I_w) + \theta_h$. If he chooses not to, $U_h = yI_h + (1-x)I_w + \theta_h$. Then, the wife either chooses to divorce the husband or not. If the husband chooses to inflict violence and the wife does not get divorced, her utility becomes $U_w = (1-y)I_h + (x-\tau)I_w + \theta_w - v_w$ where $v_w > 0$. In the absence of violence, her utility becomes $U_w = (1-y)I_h + xI_w + \theta_w$ and she stays with the husband assuming $(1-y)I_h + xI_w + \theta_w > S_w$. The husband's utility is high enough to stay with the wife in the absence of violence. When they are divorced, the wife and the husband receive S_w and S_h, respectively, their individual labor market potential. To highlight again, the level of satisfaction with the marriage remains private information for each spouse throughout the marriage.

Decisions. The solution of the game between the husband and wife can be found via backward induction. The wife chooses between staying in the marriage and getting divorced from the husband, given the husband's decision to choose violence. She will leave her husband if her utility of staying in a violent marriage is smaller than the utility of being divorced, as in

$$\underbrace{(1-y)I_h + (x-\tau)I_w + \theta_w - v_w}_{\text{Utility of being in a violent marriage}} \leq \underbrace{I_w}_{\text{Utility of being single}} . \tag{3}$$

Based on equation (3), there is a threshold value of the wife's private level of satisfaction with the marriage, $\overline{\theta_w}$, that makes her indifferent between remaining together with and being divorced from her husband. The value of $\overline{\theta_w}$ is

$$\overline{\theta_w} = (1-x+\tau)I_w - (1-y)I_h + v_w. \tag{4}$$

Recall that the cumulative distribution function of θ_w is given by F_w. Therefore, the probability of the wife divorcing her husband is $P(\theta_w \leq \overline{\theta_w}) = F_w(\overline{\theta_w})$. Observe that as I_w increases, the probability of divorce increases. This is because an increase in the wife's personal earnings is improving her outside option. Also, when I_h increases, the probability of separation decreases. This is because the wife enjoys a portion of the husband's personal earnings if she remains married.

Since the level of satisfaction with the marriage is private information, the husband chooses to inflict violence only knowing the probability that she will initiate divorce if her utility of staying together is less than her utility of being divorced. He compares the expected utility of choosing violence, EU_h, with utility of not choosing to be violent. His expected utility of inflicting violence on his wife is

$$EU_h = F_w(\overline{\theta_w})I_h + [1 - F_w(\overline{\theta_w})][yI_h + (1-x+\tau)I_w + \alpha(I) + \theta_h]. \tag{5}$$

The husband chooses to be violent if

$$\underbrace{EU_h}_{\text{Expected utility of inflicting violence}} \geq \underbrace{yI_h + (1-x)I_w + \theta_h}_{\text{Utility from not inflicting violence}} . \tag{6}$$

Based on equation (6), there is a threshold value of the husband's private level of satisfaction with the marriage, $\overline{\theta_h}$, that makes the husband indifferent between choosing violence (to extract a portion of the wife's resources) and not being violent. The value of $\overline{\theta_h}$ is determined by

$$\overline{\theta_h} = (1-y)I_h - (1-x)I_w + \frac{1-F_w(\overline{\theta_w})}{F_w(\overline{\theta_w})}[\tau I_w + \alpha(I)]. \tag{7}$$

Since violence increases the risk of divorce, a husband who is highly satisfied with the marriage (high θ_h) has more to lose and is less likely to be violent. However, if the perceived gains from violence ($\tau I_w + \alpha(I)$) are large enough, he may still choose violence despite the risk.

The probability of the husband choosing violence is $P(\theta_h < \overline{\theta_h}) = F_h(\overline{\theta_h})$. In the last component of equation (7), let $R(I_w, I_h, \tau) = \frac{1-F_w(\overline{\theta_w})}{F_w(\overline{\theta_w})}$. R is the ratio of the probability of remaining married divided by the probability of being divorced (the odds of remaining married). R is decreasing in I_w, ρ and τ and increasing in I_h.[5] As I_w increases, the woman's odds of remaining married will be lower since she has a better option.[6] The last component of equation (7) shows that in a marriage, the husband gains utility from violence via extracting monetary resources from the wife, τI_w, and non-monetary benefit from violence, $\alpha(I)$. The last component of the equation also shows how the wife's productivity in the labor market affects the husband. As the wife's productivity in the labor market increases, I_w increases. The husband's utility from extracting his wife's monetary resources via violence will be higher and his non-monetary benefit from inflicting violence will be lower.

Proposition. *For a high enough I_w and ρ or small enough τ, an increase in I_w decreases domestic violence. For a large enough income effect ($|\alpha'|$), any increase in earnings (the wife's or the husband's) decreases domestic violence.*
Proof. See Appendix A5.3.

When the wife's outside option due to her job, I_w, is high enough, the husband's extraction rate, τ, and marginal non-monetary benefit from violence with the wife's contribution to household earnings, $\frac{\partial \alpha(I)}{\partial I_w}$, is small enough, violence decreases. Moreover, by keeping the wife's earnings constant, an increase in the husband's earnings may also reduce violence. This happens when his

[5] See Appendix A5.3 for proof.

[6] On the contrary, as I_h increases, the odds of remaining married will be higher since she receives utility from the husband's income within the marriage.

marginal non-monetary benefit from violence with his contribution to household earnings $\frac{\partial \alpha(I)}{\partial I_h}$, is small enough. Importantly, the proposition also highlights that violence decreases when the husband's cost of incapacitating his wife's is high enough.

Based on the theoretical framework, mill exposure can decrease domestic violence via an increase in women's outside options, an increase in household resources (either via an increase in the wife's and/or husband's earnings) and an increase in the cost of the wife's incapacitation for the husband. Although I did not incorporate exposure reduction (decline in the time couples spend together) into the model, it is also a plausible mechanism.

A5.3 Proofs

The probability of the husband choosing violence is

$$F_h(\overline{\theta_h}) = F_h\Big[(1-y)I_h - (1-x)I_w + R(I_w, I_h, \tau)[\tau I_w + \alpha(I)]\Big] \tag{8}$$

where $R(I_w, I_h, \tau) = \frac{1-F_w(\overline{\theta_w})}{F_w(\overline{\theta_w})} = (F_w[(1-x+\tau)I_w - (1-y)I_h + v)])^{-1} - 1$. R is the ratio of the probability of remaining married divided by the probability of being separated (odds of remaining married). R is decreasing in I_w and τ and increasing in I_h as shown below.

$$\frac{\partial R(I_w, I_h, \tau)}{\partial I_w} = -F_w(.)^{-2} f_w(.)[1-x+\tau]\rho < 0. \tag{9}$$

$$\frac{\partial R(I_w, I_h, \tau)}{\partial I_h} = -F_w(.)^{-2} f_w(.)[-(1-y)] > 0. \tag{10}$$

$$\frac{\partial R(I_w, I_h, \tau)}{\partial \tau} = -F_w(.)^{-2} f_w(.) I_w < 0. \tag{11}$$

For simplicity, I took the derivatives as if I_w is a constant since the sign of the derivatives won't change even if I take I_w as $I_w(\rho)$ based on my assumptions and chain rule. Suppose $I_w(\rho) = I_w\rho$. Then $\frac{\partial R(I_w, I_h, \tau)}{\partial \rho} = -F_w(.)^{-2} f_w(.)[1-x+\tau] I_w < 0$. Thus, R is also decreasing in ρ.

I do comparative statics on $\overline{\theta_h}$ below to capture how the probability of the husband choosing violence, $F_h(\overline{\theta_h})$, changes with respect to an increase in I_w and I_h.

A5.3.1 Increase in I_w

For simplicity, first, I took the derivatives as if I_w is a constant since the sign of the derivatives won't change even if I take I_w as $I_w(\rho)$ based on my assumptions and chain rule. I will relax this at the end of the comparative statics.

$$\frac{\partial \overline{\theta_h}}{\partial I_w} = -(1-x) + \frac{\partial R(I_w, I_h, \tau)}{\partial I_w}[\tau I_w + \alpha(I)] + \Big[\tau + \frac{\partial \alpha(I)}{\partial I_w}\Big] R(I_w, I_h, \tau) \tag{12}$$

$$= -(+) + (-)[+] + [(+) + (-)](+) \tag{13}$$

The derivative cannot be signed (even if $\alpha(I) = \frac{\partial \alpha(I)}{\partial I_w} = 0$, in the absence of an income effect). The second derivative below also cannot be signed and thus $\frac{\partial \overline{\theta_h}}{\partial I_w}$ is not monotonic in I_w.

$$\frac{\partial^2 \overline{\theta_h}}{\partial I_w^2} = \frac{\partial^2 R(I_w, I_h, \tau)}{\partial I_w^2}[\tau I_w + \alpha(I)] + 2\tau \frac{\partial R(I_w, I_h, \tau)}{\partial I_w},$$

where $\frac{\partial^2 R(I_w, I_h, \tau)}{\partial I_w^2} = 2F_w(.)^{-3} f_w(.)(1-x+\tau) f_w(.)[1-x+\tau] + f'_w(.)[1-x+\tau]^2[-F_w(.)^{-2}]$.

The limit of $\frac{\partial \overline{\theta_h}}{\partial I_w}$ when I_w goes to infinity is as follows.

$$\lim_{I_w \to +\infty} \frac{\partial \overline{\theta_h}}{\partial I_w} = -(1-x) + \lim_{I_w \to +\infty} \frac{\partial R(I_w, I_h, \tau)}{\partial I_w}[\tau I_w + \alpha(I)] + \lim_{I_w \to +\infty} \Big[\tau + \frac{\partial \alpha(I)}{\partial I_w}\Big] R(I_w, I_h, \tau). \tag{14}$$

$\lim_{I_w \to +\infty} \frac{\partial \overline{\theta_h}}{\partial I_w} = -\infty$ because $\frac{\partial R(I_w, I_h, \tau)}{\partial I_w} < 0$ and $\lim_{I_w \to +\infty} R(I_w, I_h, \tau) = \lim_{I_w \to +\infty} \frac{1 - F_w(\overline{\theta_w})}{F_w(\overline{\theta_w})} = 0$. Thus, $\frac{\partial \overline{\theta_h}}{\partial I_w} < 0$ when I_w is high enough.

$\frac{\partial \overline{\theta_h}}{\partial I_w}$ is also not monotonic in τ and $\alpha(I)$ since $\frac{\partial^2 \overline{\theta_h}}{\partial I_w \partial \tau} = \frac{\partial^2 R(I_w, I_h, \tau)}{\partial I_w \partial \tau}[\tau I_w + \alpha(I)] + I_w \frac{\partial \kappa(I_w, I_h, \tau)}{\partial I_w} + R(I_w, I_h, \tau) + \frac{\partial R(I_w, I_h, \tau)}{\partial \tau}\Big[\tau + \frac{\partial \alpha(I)}{\partial I_w}\Big]$ and $\frac{\partial^2 \overline{\theta_h}}{\partial I_w \partial \alpha(I)} = \frac{\partial R(I_w, I_h, \tau)}{\partial I_w} + R(I_w, I_h, \tau)$. $\frac{\partial \overline{\theta_h}}{\partial I_w}$ is also negative when τ is small enough and $\frac{\partial \alpha(I)}{\partial I_w}$ is small enough (income effect is large enough).

$$\lim_{\tau \to 0} \frac{\partial \overline{\theta_h}}{\partial I_w} = -(1-x) + \lim_{\tau \to 0} \frac{\partial R(I_w, I_h, \tau)}{\partial I_w}[\tau I_w + \alpha(I)] + \lim_{\tau \to 0} \Big[\tau + \frac{\partial \alpha(I)}{\partial I_w}\Big] R(I_w, I_h, \tau) < 0. \tag{15}$$

$$\lim_{\frac{\partial \alpha(I)}{\partial I_w} \to -\infty} \frac{\partial \overline{\theta_h}}{\partial I_w} = -(1-x) + \lim_{\frac{\partial \alpha(I)}{\partial I_w} \to -\infty} \frac{\partial R(I_w, I_h, \tau)}{\partial I_w}[\tau I_w + \alpha(I)] + \lim_{\frac{\partial \alpha(I)}{\partial I_w} \to -\infty} \Big[\tau + \frac{\partial \alpha(I)}{\partial I_w}\Big] R(I_w, I_h, \tau) < 0. \tag{16}$$

Suppose I_w is not a constant and $I_w(\rho) = I_w \rho$. $\frac{\partial \overline{\theta_h}}{\partial I_w}$ is also negative when ρ is high enough and if $\tau \to 0$ or $I_w \to +\infty$ or $\frac{\partial \alpha(I)}{\partial I_w} \to -\infty$. This captures that when the husband's cost of incapacitating his wife is high enough, he is less likely to inflict violence.

$$\lim_{\rho \to 1} \frac{\partial \overline{\theta_h}}{\partial I_w} = -(1-x)\rho + \lim_{\rho \to 1} \frac{\partial R(I_w, I_h, \tau)}{\partial I_w}[\tau I_w \rho + \alpha(I)] + \lim_{\rho \to 1} \Big[\tau \rho + \frac{\partial \alpha(I)}{\partial I_w}\Big] R(I_w, I_h, \tau) < 0 \tag{17}$$

if $\tau \to 0$ or $I_w \to +\infty$ or $\frac{\partial \alpha(I)}{\partial I_w} \to -\infty$.

The first term, $-(1-x)I_w$, shows that when the wife's productivity ρ increases, incapacitating her becomes more costly for the husband because her earnings I_w rise. This reinforces the idea that high-productivity women are less likely to experience violence, as their incapacitation would significantly reduce household resources.

Thus, even if $\alpha(I) < 0$ and the income effect ($\frac{\partial \alpha(I)}{\partial I_w}$) were zero, the incapacitation cost mechanism would still operate as long as the wife's productivity remains sufficiently high, making violence costly for the husband.

Overall, my goal with this simple theoretical framework is not to impose one mechanism over another but to incorporate multiple mechanisms studied in the literature—such as the effect of income, extractive violence, and changes in women's outside options—while also introducing the incapacitation cost mechanism, which has not been highlighted in prior research.

A5.3.2 Increase in I_h

$$\frac{\partial \overline{\theta_h}}{\partial I_h} = (1-y) + \frac{\partial R(I_w, I_h, \tau)}{\partial I_h}[\tau I_w + \alpha(I)] + \Big[\frac{\partial \alpha(I)}{\partial I_h}\Big] R(I_w, I_h, \tau) \tag{18}$$

$$= (+) + (+)[+] + [-](+) \tag{19}$$

If $\alpha(I) = \frac{\partial \alpha(I)}{\partial I_h} = 0$ (in the absence of an income effect), $\frac{\partial \overline{\theta_h}}{\partial I_h} > 0$. An increase in husband's earnings increases violence in the household. This is because now the increase in household resources channel (e.g. stress relief) is no longer playing a role in violence and an increase in his earnings increases the benefit of marriage for the wife. Thus, she will less likely to leave when she experiences domestic violence. These make the husband more likely to inflict violence on his wife. $\frac{\partial \overline{\theta_h}}{\partial I_h} < 0$ when $\frac{\partial \alpha(I)}{\partial I_h}$ is small enough (the income effect is large enough) as in

$$\lim_{\frac{\partial \alpha(I)}{\partial I_h} \to -\infty} \frac{\partial \overline{\theta_h}}{\partial I_h} = (1-y) + \lim_{\frac{\partial \alpha(I)}{\partial I_h} \to -\infty} \frac{\partial R(I_w, I_h, \tau)}{\partial I_h}[\tau I_w + \alpha(I)] + \lim_{\frac{\partial \alpha(I)}{\partial I_h} \to -\infty} \Big[\frac{\partial \alpha(I)}{\partial I_h}\Big] R(I_w, I_h, \tau) < 0. \tag{20}$$