Polluting Industries and Agricultural Productivity: Evidence from Mining in Ghana

Fernando M. Aragón* Juan Pablo Rud†

August 2014

Abstract

This paper examines the effect of polluting industries on agricultural productivity. The focus is on large-scale gold mining in Ghana which, similar to other fuel intensive activities, releases environmental pollutants with the potential to have negative effects on crop health and key agricultural inputs. Guided by a consumer-producer household framework, we estimate an agricultural production function that incorporates the effects of pollution. We find that farmers located near gold mines experienced a reduction in total factor productivity of almost 40% between 1997 and 2005, relative to those farther away. Consistent with this result, we document higher concentrations of air pollutants and an increase in rural poverty near mines. We also explore whether mining could be affecting agricultural productivity in other ways, such as by reallocating workers or inducing changes in agricultural practices. However, we find no evidence supporting these alternative channels. Our results highlight an important externality, namely losses in agricultural productivity, through which polluting industries can affect living conditions in rural areas.

*Department of Economics, Simon Fraser University, Burnaby, British Columbia, V5A 1S6, Canada; Tel: +1 778 782 9107; Email: faragons@sfu.ca

†Department of Economics, Royal Holloway, University of London, Egham, Surrey, TW20 0EX, United Kingdom; Tel: +44 (0)1784 27 6392; Email: juan.rud@rhul.ac.uk

1 Introduction

The investigation of the social costs of environmental pollution has mostly focused on its negative effects on human health.[1] Recent studies find that pollution, through its effect on health, can also hinder human capital accumulation (Currie et al., 2009), labour supply (Hanna and Oliva, 2011) and labour productivity (Graff Zivin and Neidell, 2013). However, the economic literature has paid less attention to the effects of pollution on other economic outcomes, such as agricultural productivity. This is surprising given the existing biological evidence linking pollution to reductions in crop health and yields (Heck et al., 1982; Miller, 1988; Marshall et al., 1997) and degradation of key agricultural inputs, such as water and soil (Menz and Seip, 2004; U.S. Environmental Protection Agency, 2012).

This paper addresses this gap in the economic literature by examining how polluting industries affect agricultural productivity in a context where traditional farming is the main source of livelihood. Quantifying this externality is important to inform the debate on environmental policies and to assess the net benefits of (potentially) polluting activities, such as urban growth and extractive industries, which may occur in the vicinity of agricultural areas.

We study gold mining in Ghana as it presents three useful features for our purposes. First, most gold production is done in large-scale, modern mines. These mines are heavily mechanised and release air pollutants similar to other fuel-intensive activities, such as power plants and urban traffic. These pollutants can be carried over long distances and, in high concentrations, can build up in the environment and have cumulative effects.[2] Second, large gold mines have little interaction with local economies: they hire few local workers, buy few local products, and almost none of its profits are locally distributed.[3] This effectively shuts down a number of alternative channels through which mining activity can affect agricultural activities. Finally, gold mines in Ghana are located in the vicinity of fertile agricultural lands where important cash crops, such as cocoa, are cultivated.

We use micro-data from household surveys with information on agricultural practices for the years 1997 and 2005 and detailed data on location of gold mines and households. To study the

[1]See Graff Zivin and Neidell (2013) and Currie et al. (2013) for a comprehensive review of this literature.

[2]Gold mining can generate other industry-specific stock pollutants, e.g. cyanide spills and acidic discharges. These pollutants are mostly carried by water or localised in the close vicinity of mine sites.

[3]Modern mining is often associated with this type of 'enclave effect'. See Aragon and Rud (2013) for a discussion. Anecdotal evidence for Ghana can be found in Aryeetey et al. (2007)

effect of pollution on agricultural productivity, we estimate an agricultural production function augmented with pollution effects. This allows us to examine how total factor productivity, the residual output conditional on observable inputs, is affected by exposure to mining-related pollution. We use cumulative gold production to proxy for the stock of these pollutants.

A main empirical challenge is that agricultural productivity may be systematically different in both mining and non-mining areas. To overcome this concern, we use a difference-in-difference approach exploiting two sources of variation: distance from households to the nearest mine and changes in mining production. The main identification assumption is that the change in agricultural productivity over time in both areas would be similar in the absence of mining.

A second challenge is the endogeneity of input use in estimating agricultural production functions, which has long been recognised in the empirical literature.[4] However, due to data limitations, we are unable to implement the standard solutions. Instead, we use the analytical framework of consumer-producer households (Benjamin, 1992; Bardhan and Udry, 1999) to derive an empirical strategy. We show that, in the presence of imperfect input markets, endowments are a good predictor of input use. Consequently, we use farmers' input endowments, such as land holdings and household size, as instruments. This instrumental variable strategy exactly identifies our production function parameters if the instruments are uncorrelated with unobserved productivity shifters. We also investigate the case where there is some correlation between instruments and unobserved heterogeneity by using the partial identification strategy of Nevo and Rosen (2012). We find that our results are robust to small correlations of this type.

We find evidence of a significant reduction in agricultural output and total factor productivity attributed to mining activities. Our estimates suggest that an increase of one standard deviation in our measure of gold production is associated with a 10 percent decline in productivity in areas within 20 km of a mine. Given the increase in mining activity between 1997 and 2005 this implies that the average agricultural productivity in farms in the vicinity of mines decreased around 40% relative to areas farther away. Similar results are obtained when using partial measures of productivity such as crop yields. The results are robust to alternative estimation methods and model specifications, and are driven by proximity to operating mines.[5] An important implication of the consumer-producer framework is that a reduction in agricultural

[4]See Ackerberg et al. (2006) and references therein for a discussion of alternative methods.

[5]A placebo test shows no changes on productivity of farmers close to new mining projects that were not operating in the period of analysis.

production affects directly a household's consumption possibilities. Indeed, we find that poverty in mining areas shows a relative increase of around 18 percent.

Having established that mining is associated with a reduction in farmers' productivity, we look for evidence of pollution. Using satellite imagery we find that the concentration of nitrogen dioxide (NO_2), a key indicator of air pollution, is higher in locations where mines operate and declines with distance. We cannot test directly whether pollution reduces labour productivity, the quality of soils, or the health of plants. However, we provide suggestive evidence that the effect is not entirely driven by a reduction of labour productivity. As an example, a back-of-the envelope calculation using the structural estimates suggests that the reduction of labour productivity would need to be very large (around 80%) to fully account for the observed drop in total factor productivity.[6]

Finally, we also investigate alternative channels that could explain the reduction in productivity. In particular, we focus on differences in the composition of agricultural workers, e.g. due to selective migration or reallocations towards non-agricultural activities.[7] We also look for changes in agricultural practices and investments that might result from a weakening of property rights, in areas where mining licences are granted.[8] However, we do not find any evidence of changes in observable characteristics of agricultural workers, in workers' occupation or in agricultural practices that are consistent with the lower productivity we observe near mining areas.

In addition to the aforementioned environmental economics literature studying the impacts of pollution, this paper also contributes to a growing literature studying the local impact of natural resources.[9] Our contribution is to quantify the potential costs, in terms of agricultural productivity and rural income, associated to pollution from extractive industries and highlight a dimension that is currently absent in the policy debate. This omission may overestimate the contribution of extractive industries to local economies and lead to insufficient compensation and mitigation policies.

[6] To put this figure in context Graff Zivin and Neidell (2012) find that a decrease on ozone of 10 parts per billion (ppb) increases worker's productivity by 5.5%. In their study, the average ambient ozone is under 50 ppb with a standard deviation of 13 ppb.

[7] This re-allocation of resources may occur if mines hire local workers or create a local demand boom, as in Aragon and Rud (2013)

[8] Besley (1995) shows, in the context of rural Ghana, that investments in crops with high return in the long run, such as cocoa trees, are lower when property rights are less secure.

[9] See, for example, Caselli and Michaels (2013) for (negative) political economy channels, Aragon and Rud (2013) for positive market channels, and Kotsadam and Tolonen (2013) for a gender-specific reallocation of labour.

The rest of the paper is organised as follows. Section 2 briefly summarises the relationship between pollution and agricultural productivity, and describes the case of gold mining in Ghana. Section 3 provides a conceptual framework, describing the data and empirical strategy. Section 4 presents the results and discusses several possible challenges to our empirical strategy and to the interpretation of our results. Section 5 explores the effect on poverty, while Section 6 concludes.

2 Background

2.1 Mining and pollution

Modern mining technologies have the potential to pollute the environment in several ways. First, significant amounts of air pollutants may be generated through the use of heavy machinery, smelters and refineries and from blasting operations.[10] At low concentrations, air pollutants are short lived: they are dissipated or absorbed by the environment. However, if emissions are relatively large, they can be carried away over long distances and can be directly absorbed by plants or deposit on the ground as acid rain.[11]

Second, mines can also generate industry-specific pollutants, such as cyanide, heavy metals, or acid mine drainage (Salomons, 1995; Dudka and Adriano, 1997). Cyanide, for example, is generally reprocessed but there is the risk of leakages during transportation or seeping from dumping tailings. Acid mine drainage occurs when sulphide minerals are exposed. Combined with air and water, they form a very acidic effluent. Importantly for our analysis, these pollutants are mostly carried by surface water. This may limit the pollutants' impact on agriculture in the Ghanaian case, where most crops are rainfed. For this reason, in the rest of the paper we focus on air pollutants. In Section 4.2, however, we also explore the role of pollutants carried by surface waters.

[10] These air pollutants include nitrogen oxides (NO_x, namely NO and NO_2), sulphur dioxide (SO_2), ozone (O_3) and particulate matter

[11] Acid rain is formed when emissions of NO_x or SO_2 react with water in the atmosphere to produce acids. It contributes to soil degradation and can have cumulative negative effects (Menz and Seip, 2004).

2.2 Pollution and agricultural productivity

Air pollution has been documented to affect agricultural productivity in at least three ways. First, evidence in biological sciences (Heck et al., 1982; Miller, 1988; Marshall et al., 1997) suggests that air pollutants, such as nitrogen oxides, have a sizeable negative effect on crop yields. For example, Emberson et al. (2001), Maggs et al. (1995), and Marshall et al. (1997) find reductions of around 20 to 60 percent in the yield of crops such as rice, wheat, and beans that are exposed to polluted air from urban centres located as far as 15 km away.[12] Second, pollution can generate acid rain that deteriorates soil quality, by changing its chemistry or reducing the concentration of important plant nutrients. These effects are cumulative and long-lived.[13] Finally, recent studies find evidence of a negative impact of air pollution on labour supply and productivity (Graff Zivin and Neidell, 2012; Hanna and Oliva, 2011), mostly due to its effect on human health.

2.3 Mining in Ghana

Our empirical analysis pertains to the case of gold mining in Ghana. Most of the gold (around 97%) is extracted by modern, large-scale mines located in three regions: Western, Ashanti and Central.[14] These mines, similar to other modern mines in the world, are capital intensive, highly mechanised operations. They are located in rural areas, amidst fertile agricultural land, and have little interaction with local economies: they hire few local workers, buy few local products, their profits are not distributed among local residents, and only a small fraction of the fiscal revenue is allocated to local authorities (Aryeetey et al., 2007).

Due to data availability, we focus on two years: 1997 and 2005. As shown in Figure 1, before 1997 gold production was increasing from low levels of production. This was mostly driven by the expansion of one mine, Obuasi.[15] After 1997, gold production plateaus, but at a higher level. Table 1 shows that the aggregate cumulative production has almost tripled between 1997

[12] Most of the available evidence comes from developed countries. The above mentioned studies, however, document the effect of pollution in developing countries such as India, Pakistan and Mexico.

[13] For a summary of this evidence see websites of the U.S. and Canada environmental agencies (`http://www.epa.gov/acidrain/effects/forests.html` and `http://www.ec.gc.ca/air/default.asp?lang=En&n=7E5E9F00-1ws0EF0FB73`).

[14] The rest is produced by small artisanal operations that are usually owned by locals and by informal miners called *galamseys*. These use a similar labour-intensive, small-scale technology.

[15] The main results are robust to excluding observations in the vicinity of Obuasi mine (see Table A.8 in the online appendix.

and 2005 and that there is substantial variation across mines. Extraction occurs at a greater number of locations and many of these operations, such as Tarkwa, Bibiani and Damang, were new or experienced a significant expansion. We exploit these differences in gold production by mine in our empirical analysis.

Figure 1: Total gold production (in Tonnes), by year

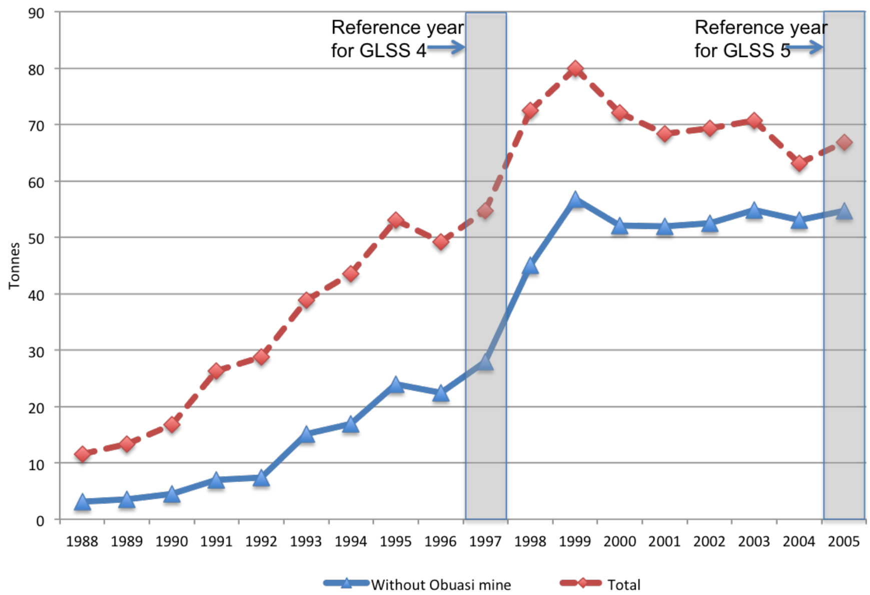

Source: U.S. Geological Service, *The Mineral Industry of Ghana* 1994-2004, Infomine, and Aryeetey et al. (2007).

There are no systematic data on the concentration of pollutants in the vicinity of mining sites, even though some case studies in mining areas report the presence of heavy metal pollutants and levels of particulate matter above international admissible levels.[16] The levels of concentration decay as distance to mining sites increases, probably due to air dispersion (see for example, Armah et al. (2010) and Tetteh et al. (2010)). As these case studies do not cover all relevant areas and years, they are unsuitable for our analysis. Instead, we use mines' cumulative gold production over the relevant period as a proxy for the generation of pollutants that accumulate in the environment over time.[17]

[16]Only since 2009 Ghana's Environmental Protection Agency (EPA) has started assessing, and reporting, the environmental compliance of mines (see `http://www.epaghanaakoben.org/`). Of the 9 operative gold mines studied, 7 were red-flagged as failing to comply environmental standards. These mines were considered to pose serious risks due to toxic and hazardous waste mismanagements and discharge.

[17]These pollutants are called *stock* pollutants. In contrast, *flow* pollutants are dissipated or absorbed by the environment.

Table 1: Cumulative gold production by mine, in Tonnes

Mine name	Type	Cumulative production		
		Up to 1997	Up to 2005	Diff.
Bibiani	open pit	0.0	51.2	51.2
Bogoso/Prestea	open pit, underground and and tailings	23.9	55.9	32.0
Central Ashanti	open pit	5.4	9.7	4.3
Damang	open pit	0.0	73.6	73.6
Dunkwa placer	placer	1.2	1.2	0.0
Essase placer	placer	2.8	12.4	9.6
Iduapriem/Teberebie	open pit	19.6	61.2	41.6
Konong/Obenamasi	open pit	1.5	1.5	0.0
Obotan	open pit	2.2	19.4	17.3
Obuasi	open pit and underground	204.3	346.3	142.0
Tarkwa	open pit and underground	9.4	121.0	111.6
Wassa	open pit	0.0	10.3	10.3
TOTAL		270.3	763.7	493.4

Note: Cumulative production is calculated adding annual production from year 1988 to 1997, and from 1988 to 2005, respectively. Data collected from U.S. Geological Service, *The Mineral Industry of Ghana* 1991-2004, Infomine, and Aryeetey et al. (2007).

3 Methods

3.1 A consumer-producer household

In this section we lay down a simple analytical framework to guide the empirical analysis. In particular, we extend a standard model of consumer-producer households (Benjamin, 1992; Bardhan and Udry, 1999) to understand how an expansion of mining activities can generate adjustments in the optimal behaviour of households.

We assume that households (farmers) are both consumers and producers of an agricultural good with price $p = 1$. Households have an idiosyncratic productivity A and use labour (L) and land (M) to produce the agricultural good $Q = F(A, L, M)$, where F is a concave production function. Households have endowments of labour and land (E^L, E^M). They can use these endowments as inputs in their farms, sell them in local input markets (L^s, M^s) at prices w and r, or, in the case of labour, consume it as leisure. As producers, households can buy additional labour and land (L^b, M^b).

Households maximise utility $U(c, l)$ over consumption c and leisure l, subject to the budget constraint $c = F(A, L, M) - w(L^b - L^s) - r(M^b - M^s)$, and the endowment constraints $L =$

$E^L + L^b - L^s - l$ and $M = E^M + M^b - M^s$.

We assume households are heterogeneous in their access to markets for inputs. In particular, there are two types of farmers: (1) unconstrained farmers, who operate as in perfectly competitive input markets, and (2) fully-constrained farmers, who can neither buy nor sell inputs.[18] The assumption of imperfect input markets is reasonable in the context of weak property rights of rural Ghana.[19] [20]

In the case of unconstrained farmers, the maximization problem follows the separation property: the household chooses the optimal amount of inputs to maximise profits and, separately, chooses consumption and leisure levels, given the optimal profit. From standard procedures, the optimal levels of inputs and output, $L^*(A, w, r)$, $M^*(A, w, r)$ and $Q^*(A, w, r)$, depend only on total factor productivity and input prices.

In the case of fully-constrained farmers the optimal input decisions are shaped by their endowments. Since the opportunity cost of land is zero, they will use all their land endowment, $M^* = E^M$. However, in the case of labour farmers face a trade-off between leisure and income. Solving the household's problem, the optimal level of labour, $L^*(A, E^L, E^M)$, is a function of total factor productivity and input endowments.[21]

In this framework, there are two possible channels for mining to affect agricultural output and household consumption. First, mines could increase demand for local inputs (input competition). This may lead to an increase in input prices and, through that channel, reduce input use and agricultural output among unconstrained farmers. Similar effects would occur if, for example, mines reduce supply of inputs due to land grabbings or population displacement. There would be, however, no effect on productivity A.[22]

Second, mining-related pollution may affect crop health and yields as well as the quality of

[18] Results would not change qualitatively if we allow for partially constrained farmers.

[19] Data show that, in the area of study, input markets are thin: around 8% of available land is rented, and only 1.4% of the total farm labour (in number of hours) is hired. As shown in Table A.2 in the online appendix, endowments are a very strong predictor of input use.

[20] Besley (1995), for example, documents the co-existence of traditional and modern property right systems in West Ghana. Some farmers have limited rights to transfer property of land, and in many cases require approval from the community while others do not face this constraint. Botchway (1998) also discusses the customary framework that rules the right to trade land in Ghana. Similar arguments can be made about labour markets, due to market incompleteness, farmers' preference for working on their own land, or imperfect substitutability of household and hired labour.

[21] For a fully constrained farmer, the household's problems simplifies to $\max U(c, l)$ subject to $c = F(A, L, E^M)$ and $L = E^L - l$. The first order condition is $U_c F_L = U_l$.

[22] This remark depends, however, on the assumption that input type does not change. We explore the validity of this assumption in the empirical analysis.

inputs, as discussed above. This would imply a reduction in output even if the quantity of inputs used remains unchanged. In terms of the model, this represents a drop in productivity A. This would have an unambiguous negative effect on agricultural output and household consumption. Additionally, it might reduce input use. Labour use might fall either by reducing labour demand for unconstrained farmers or through a substitution of labour towards leisure for constrained farmers. In the case of land, only unconstrained farmers would reduce their land use.

These results highlight the importance of studying total factor productivity to assess the effect of mining-related pollution. Other outcomes, such as agricultural output or input use, might not be very informative about the channels at play. However, this also raises an empirical challenge: unobserved heterogeneity in A can also affect input use and compromises the econometric identification of total factor productivity. In our empirical approach we rely on the model prediction that, in the presence of imperfect input markets, household endowments can be a key determinant of input use to consistently estimate production function parameters.

3.2 Empirical implementation

The aim of the empirical analysis is to explore the importance of mining-related pollution on agricultural activity. To do so, our main approach is to estimate the production function and evaluate the effect of mining on total factor productivity A.

We start by assuming the following agricultural production function:[23]

$$Y_{ivt} = A_{ivt} M_{it}^{\alpha} L_{it}^{\beta} e^{\epsilon_{it}}, \tag{1}$$

where Y is actual output, A is total factor productivity, M and L are land and labour, and ϵ_{it} captures unanticipated shocks, which is by definition uncorrelated with input decisions. All these variables vary for farmer i in locality v at time t. Other inputs, such as capital and materials (e.g. fertilisers, insecticides), are not widely used and thus excluded from the empirical analysis.[24] Their inclusion, however, does not change any of the results.

We assume that A is composed of three factors: farmers' heterogeneity (η_i), time-invariant

[23]We assume a Cobb-Douglas technology for simplicity. We also check the robustness of the results to using a, more general, CES production function (see Section B in the on-line appendix).

[24]For example, the value of tools and other capital goods is, on average, less than 1% of total output and the value of manure, seeds, fertilisers and insecticides account for less than 5%.

local economic and environmental conditions (ρ_v) and time-varying factors, potentially related to the presence of local mining activity (S_{vt}). In particular, $A_{ivt} = \exp(\eta_i + \rho_v + \gamma S_{vt})$. Note that if mining affects input availability or prices (through an input competition channel), it will change input use but would not affect productivity A so $\gamma = 0$. In contrast, if the pollution mechanism is at play, we should observe $\gamma < 0$.

As the empirical counterpart of S_{vt}, we use cumulative gold production near a farmer's locality.[25] This variable would be a reasonable proxy for exposure to pollutants under the assumption that pollutants have a cumulative effect, i.e. they are stock pollutants. As we discuss in Section 2, several pollutants released by mining operations, such as NO_2, SO_2 and heavy metals, can have negative cumulative effects on vegetation through acid rain and soil degradation.[26]

We can anticipate two main empirical challenges. The first is related to the fact that mining and non-mining areas may have systematic differences in productivity. This omitted variable problem may lead to endogeneity issues when estimating the coefficients of interest. To address this issue, we exploit time variation in the repeated cross section to compare the evolution of productivity in mining areas relative to non-mining areas. This approach is basically a difference-in-difference with a continuous treatment. In this case, proximity to a mine defines the treated and control group, while the intensity of the treatment is the cumulative amount of gold produced in nearby mines.[27] The validity of this approach relies on the assumption that the evolution of productivity in both areas would have been similar in the absence of mining.[28]

The second problem arises because both output and input choice can be affected by productivity, and hence may be simultaneously determined. Thus, unobserved heterogeneity in productivity would be reflected in the error term and create an endogeneity problem in the

[25] In the baseline specification, we define a mining area as localities within 20 km of a mine. For those areas, S_{vt} is equal to gold production in nearby mines from 1988 to the reference year of the household survey (i.e. 1997 for GLSS 4 and 2005 for GLSS 5). For areas farther than 20 km, $S_{vt} = 0$.

[26] In the empirical analysis, we also check the robustness of the results to measures of flow pollutants, i.e. short-lived pollutants, using annual gold production (see Table 5).

[27] We also use a simpler specification replacing S_{vt} by $mining\_area_v \times T_t$ where $mining\_area_v$ is an indicator of being within 20 km of a mine and T_t is a time trend. The results using this discrete treatment are consistent with the continuous case (see Table A.3 in the online appendix).

[28] In the online appendix we explore the evolution of average agricultural output in mining and non-mining areas three years with data from GLSS 2 (1988), GLSS 4 (1997) and GLSS 5 (2005). Figure A.1 shows that the evolution of output is remarkably similar in the first period (1988-1997), when gold production is relatively low, but there is a trend change in mining areas in the period when gold production increases (1998-2005). Table A.1 formally tests the similarity of trends, and subsequent change, by regressing agricultural output on $mining\ area_v \times T_t$ for both periods. Note that the similarity of trends prior to the expansion of mining is a necessary, though not sufficient, condition for the identification assumption to be valid.

estimation of the input coefficients.

We address these concerns in several ways. First, we use observable characteristics of farmers to proxy for heterogeneity, η_i. We also include district fixed effects to capture differences in average product due to local characteristics.[29] With these modifications, and taking logs, the model we estimate becomes:

$$y_{ivdt} = \alpha m_{it} + \beta l_{it} + \gamma S_{vt} + \phi Z_i + \delta_d + \psi_t + \theta mining\_area_v + \xi_{ivt}, \quad (2)$$

where y, l and m represent the log of observed output, labour and land, respectively. Z_i is a set of farmer's controls, and S_{vt} is the cumulative gold production in the proximity of a locality. δ_d and ψ_t represent district and time fixed effects, while $mining\_area_v$ is an indicator of being within 20 km of a mine. ξ_{ivt} is an error term that includes ϵ_{it} and the unaccounted farmer and locality heterogeneity. Under the assumption that use of inputs is uncorrelated to residual unobserved heterogeneity, ξ_{ivt}, we can estimate the parameters of (2) using an OLS regression.

Second, we relax the previous identification assumption and exploit the presence of some constrained farmers. In particular, we estimate a standard IV model using endowments as instruments for input use. Recall from the model that the larger the fraction of constrained households, the greater the correlation between input use and household endowments. This approach would be valid if the correlation is strong enough and if endowments affect output only through its effect on input use, i.e. endowments are not conditionally correlated to unobserved heterogeneity, ξ_{ivt}.[30]

Finally, we consider the possibility that endowments are correlated to ξ_{ivt}. This could happen, for example, if more productive farmers have systematically larger landholdings or household size, thereby invalidating the exclusion restriction of the IV strategy. However, we can make further progress by using a partial identification strategy proposed by Nevo and Rosen (2012). This methodology uses imperfect instrumental variables (IIV) to identify the set of parameter values.[31] The approach relies on two assumptions: (i) the correlation between the instrument and the error term has the same sign as the correlation between the endogenous

[29] Districts are larger geographical areas than localities v. We cannot use locality fixed effects due to the structure of the data.

[30] The interpretation of this IV strategy would be as a local average treatment effect, since the coefficients would be identified from constrained farmers only.

[31] In contrast, the standard IV approach focuses on point identification.

variable and the error term, and (ii) the instrument is *less* correlated to the error than the endogenous variable. These (set) identification assumptions are weaker than the exogeneity assumption in the standard IV and OLS approaches.[32]

3.3 Data

Our main results use a repeated cross-section of household data from the rounds 4 and 5 of the Ghana Living Standards Survey (GLSS 4 and GLSS 5).[33] These surveys were collected by the Ghana Statistical Service (GSS) between April 1998 to March 1999, and September 2005 to August 2006, respectively. Questions on agricultural activities refer to the previous 12 month-period, therefore the surveys reflect information on agricultural input and outputs mainly for years 1997 and 2005. We use these two years as the reference years to match household data with measures of mining activity.

The survey is representative at the regional level and contains several levels of geographical information of the interviewees.[34] The finer level is the enumeration area, which roughly corresponds to villages (in rural areas) and neighbourhoods (in urban areas). For each enumeration area we obtain its geographical coordinates from the GSS.[35]

We are mainly interested on two sets of variables: measures of mining activity, and measures of agricultural inputs and output.

Mining activity Our main measure of mining activity is the cumulative production of gold in the proximity of a household, the empirical counterpart of S_{vt}.[36]

For each of the mines in Table 1, we obtain geographical coordinates of their site.[37] Using a geographical information system (ArcGIS), we identify the enumeration areas within different distance brackets of each mine site. For now, we define the enumeration areas within 20 km of

[32]We refer the reader to Nevo and Rosen (2012) for a detailed exposition of the estimation method.

[33]We also use the GLSS 2, taken in 1988/89, for evaluating pre-trends in agricultural output between mining and non-mining areas. However, we do not use this dataset in the estimation of the production function since it does not contain comparable information on input use. In addition, we do not use the GLSS 3 (1993/94) because there is not available information on the geographical location of the interviewees.

[34]The highest sub-national administrative jurisdiction level is the region, followed by the district. In 2005, there were 10 regions and 138 districts. The survey also distinguishes between urban and rural areas, as well as ecological zones (coastal, savannah and forest).

[35]The GSS does not have location of enumeration areas for the GLSS 2. In this case, we extracted the location using printed maps of enumeration areas in previous survey reports.

[36]We measure this variable in hundred of tonnes.

[37]This information comes from proprietary industry reports prepared by Infomine.

mine sites as mining areas. Finally, we assign the cumulative production of each mine to its surrounding mining area, and zero for areas farther away.

Figure 2a displays a map of Ghana with the location of active gold mines between 1988 and 2005. Note that all mines are located in three regions: Western, Ashanti and Central. In the empirical section, we restrict the sample to these regions.[38] Figure 2b focuses on these three regions and depicts the enumeration areas and a buffer of 20 km around each mine. The areas within each buffer correspond to the mining areas (treated group), while the rest correspond to the non-mining areas (comparison group).

We restrict attention to medium and large-scale gold mines. We do not consider artisanal and informal gold mines for two reasons. First, the magnitude of their operations is relatively small (they represent around 3% of total gold production). Second, there is no information on their location, though anecdotal evidence suggests they are located in the vicinity of established mines. For similar reasons, we do not consider mines of other minerals (such as diamonds, bauxite and manganese). These minerals are less important than gold in Ghana's mining output. Moreover, their mine sites are concentrated in locations that overlap with existing gold operations. For example bauxite and diamonds are mined in Awaso (south of Bibiani gold mine), while manganese is extracted at the Nsuta-Wassa mine near Tarkwa.

[38] The results, however, are robust to using a broader sample.

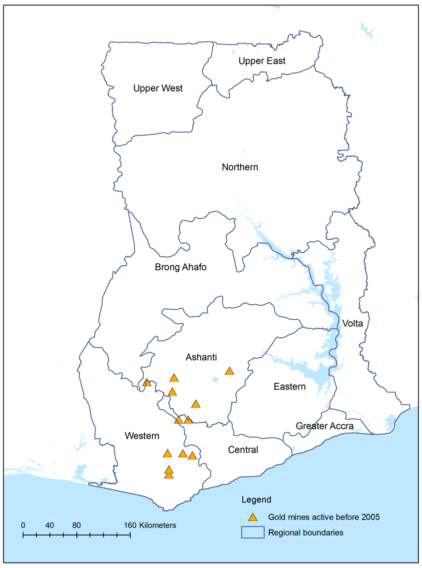

(a) Location of active gold mines

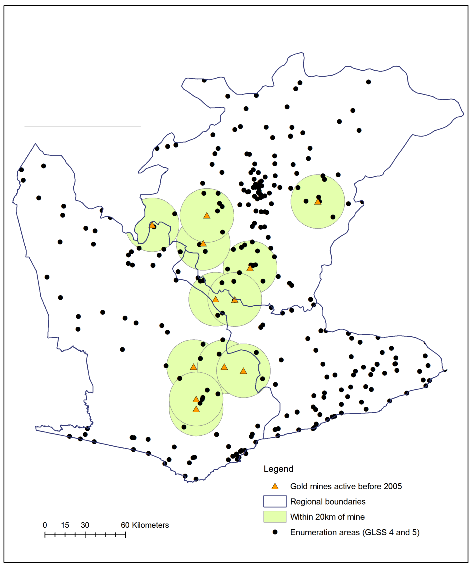

(b) Area of study and enumeration areas

Figure 2: Location of gold mines and households

Agricultural output and inputs To measure agricultural output, Y, we first obtain an estimate of the nominal value of agricultural output. To do so, we add the reported value of annual production of main crops. These category includes cash crops and staple grains such as cocoa, maize, coffee, rice, sorghum, sugar cane, beans, peanuts, etc. Then, we divide the nominal value of agricultural output by an index of agricultural prices.[39] This price index uses data from agricultural producers and varies by region and by mining and non-mining areas.[40]

We also construct estimates of the two most important agricultural inputs: land and labour. The measure of land simply adds the area of plots cultivated with major crops in the previous 12 months. To measure labour, we add the number of hired worker-days to the number of days each household member spends working in the household farm. Finally, we measure land endowment as the area of the land owned by the farmer, while the labour endowment is the number of equivalent adults in the household.

The resulting dataset contains information on agricultural inputs and output for 1,627 farmers. The farmers are located in 42 districts in three regions of south west Ghana: Western, Ashanti and Central. Table 2 presents a simplified difference-in-difference estimation of the main variables of interest, by comparing mean values in both survey rounds for farmers located in both mining and non-mining areas. A first important observation is that the log of agricultural output has shown a relative decrease near mining areas. Consistent with the consumer-producer household framework, the poverty rate in affected areas shows a relative increase. On the contrary, there is no apparent significant difference in the use of the main inputs, land and labour. There is a differential change in input prices that has the opposite sign we would expect if there were an increase in labour demand from mines. This reduction in input prices might simply reflect the lower marginal productivity of inputs due to pollution.

There are no significant differences in most farmers' characteristics, except for place of birth and land ownership. These differences, however, disappear when controlling for other farmer characteristics.

[39]The results are similar robust to using a coarser consumer price index reported by the GSS, which varies by ecological zone and by urban and rural areas (see Table A.4 in the online appendix). This consumer price has a lower geographical resolution than the one we use in this paper.

[40]In particular, we obtain data from individual farmers on unit values of cocoa and maize, the two main crops in the area of study, and their relative share in the value of agricultural output in 1997. Then, we take the median value of prices and weights by region and by mining and non-mining area, i.e., six different values every survey, and construct a Laspeyres price index.

Table 2: Mean of main variables, by GLSS and location

Variable	Within 20 km of mine		Outside 20 km of mine		Diff. columns
	GLSS 4	GLSS 5	GLSS 4	GLSS 5	(2-1) - (4-3)
	(1)	(2)	(3)	(4)	(5)
Cumul. gold prod. (MT)	41.7	84.6	-	-	-
ln(real agric. output)	6.6	6.2	6.5	6.6	-0.5***
					(0.17)
Land used (acres)	7.2	17.9	8.3	9.4	9.6
					(9.50)
Labor (days)	377.3	358.8	343.1	366.3	-41.7
					(32.00)
Land owned (acres)	11.6	21.2	12.0	13.6	8.0
					(9.65)
Nr. adults equivalents	3.6	3.4	3.9	3.5	0.2
					(0.23)
ln(relative land price)	14.4	14.1	13.9	14.1	-0.5***
					(0.10)
ln(real wage)	8.6	8.8	8.4	8.8	-0.20***
					(0.04)
Age (years)	48.0	47.9	46.6	47.4	-0.9
					(1.9)
Literate (%)	53.1	46.6	54.5	45.3	2.7
					(6.3)
Born in village (%)	45.5	60.7	54.2	41.9	27.5***
					(6.2)
Owns a farm plot (%)	69.3	88.4	54.3	83.0	-9.6*
					(5.4)
Poverty headcount (%)	15.2	26.0	33.8	17.6	27.0***
					(5.0)
Nr. Observations	162	218	551	696	

Notes: Standard errors in parentheses. * denotes significant at 10%, ** significant at 5% and *** significant at 1%. Columns 1 to 4 report mean values for the sub-sample of farmers within and outside 20 km of a mine for every round of the GLSS. Means are estimated using sample weights. By definition, cumulative production in non-mining areas is equal to zero in both periods. Column 5 displays the difference in difference of columns 1 to 4. The standard errors are in parentheses. Total number of observations is 1,627.

4 Main results

This section provides evidence that the expansion of mining activities is associated with a significant reduction in agricultural productivity. The results are robust to various specifications and estimation techniques. While unable to directly measure mining-related pollution, we use satellite imagery to show that air pollution concentrates around mining areas. We explore alternative explanations of the productivity decline, such as changes in population composition or risk of expropriation, but find no supporting evidence that these channels can explain our results. We conclude by discussing the mechanisms through which pollution could affect productivity.

4.1 Effect on agricultural productivity

Table 3 presents the main results. In column 1, we explore the relationship between agricultural output and our measure of mining activity, cumulative production in nearby mines, without controlling for input use. We note that this relationship is negative and significant, consistent with mining affecting agriculture both through pollution or through input competition, as discussed in Section 3.

To explore the likely channels driving this relation, we proceed to estimate the agricultural production function specified in equation (2). Column 2 provides OLS estimates, while column 3 estimates a 2SLS using input endowments (namely, area of land owned and the number of adults equivalents in the household) as instruments for actual input use.[41] As a reference, column 4 estimates the 2SLS regression using the interaction between a dummy of proximity to a mine and a time trend as a proxy for S_{vt}. In this case, the estimate of γ represents the average change in residual productivity in mining areas relative to non-mining areas. All regressions include a set of farmer controls, district and year fixed effects. We also use sample weights and cluster errors at district level to account for sampling design and spatial correlation of shocks.

Both approaches suggest a large negative relationship between mining and output, after controlling for input use.[42] Under the identification assumptions discussed above, we interpret this as evidence that mining has reduced agricultural productivity. This result is consistent

[41] The first stage of the 2SLS reveals a positive and significant correlation between input endowments and input use and is very strong, using standard F-test thresholds. This is consistent with the presence of imperfect input markets as discussed in Section 3.1. See Table A.2 in the online appendix for the first stage regressions.

[42] The estimates of α and β, i.e., the participation of land and labour, also seem plausible. We cannot reject the hypothesis of constant returns to scale. Using the 2SLS estimates, the p-value of the null hypothesis $\alpha + \beta = 1$ is 0.773. We obtain a similar result of constant returns to scale when estimating a CES production function.

with mining-related pollution negatively affecting agriculture.

The magnitude of the effect is economically relevant: an increase of one standard deviation in the measure of mining activity is associated with a reduction of almost 10% in agricultural productivity.[43] Using the result in column 4 and the increase in cumulative production between 1997 and 2005, the average agricultural productivity in areas proximate to mines decreased around 40% relative to areas farther away. The estimated effect on productivity is large; however, this magnitude is consistent with the biological literature that documents reductions of 30-60% in crop yields due to air pollution (see Section 2).

So far, we have assumed that areas within 20 km of mines experience most of the negative effect. Implicitly, this approach assumes that the effect of mining declines with distance. To explore this issue further, we estimate equation (2) replacing S_{vt} with a linear spline of distance to a mine, $\sum_c \gamma^d (distance_v^d \times T_t)$ where $distance_v^d = 1$ if enumeration area v is in distance bracket d, and T_t is a time trend. This specification treats distance more flexibly and allow us to compare the evolution of farmers' productivity at different distance brackets from the mine relative to farmers farther way (the comparison group is farmers beyond 50 km).

Figure 3 presents the estimates of γ^d. Note that the effect of mining on productivity is (weakly) decreasing in distance. Moreover, the loss of productivity is significant (at 10% confidence) within 20 km of mines, but becomes insignificant in farther locations. This result provides the rationale for the threshold of 20 km around mines.[44]

Columns 5 and 6 examine the effect of mining on crop yields, i.e. physical production per unit of land. This is a standard measure of agricultural productivity that abstracts from output aggregation and deflation issues. However, it is not informative about the source of changes (whether input use or A). We focus on the yields of cocoa and maize, the two most important crops in south west Ghana. In both cases, we estimate an OLS regression including farmer controls and district fixed effects and we also find a negative and significant relation between mining and productivity.[45]

As a further check, we use the imperfect instrumental variable approach developed by Nevo

[43]The average value of the measure of mining activity (cumulative gold production within 20 km in hundreds of Tonnes) increased from 0.417 in 1997 to 0.846 in 2005. The standard deviation of this variable is 0.617.

[44]Tables A.5 and A.6 in the online Appendix replicate all the main results and robustness checks using this specification. Results consistently show effects within 20km.

[45]We do not control for inputs since we do not have estimates of labour use by crop. However, including total input use does not change the results.

Table 3: Mining and agricultural productivity

	ln(real agricultural output)				ln(yield cocoa)	ln(yield maize)
	(1)	(2)	(3)	(4)	(5)	(6)
Cumulative gold prod. within 20 km.	-0.149*	-0.176**	-0.170**		-0.509*	-0.420***
	(0.085)	(0.085)	(0.085)		(0.298)	(0.103)
Within 20 km of mine × GLSS 5				-0.565**		
				(0.240)		
ln(land)		0.631***	0.676***	0.678***		
		(0.038)	(0.047)	(0.046)		
ln(labor)		0.209***	0.352***	0.346***		
		(0.033)	(0.110)	(0.109)		
Estimation	OLS	OLS	2SLS	2SLS	OLS	OLS
Observations	1,627	1,627	1,627	1,627	948	605
R-squared	0.221	0.445	0.435	0.438	0.349	0.409

Notes: Robust standard errors in parentheses. Standard errors are clustered at district level. * denotes significant at 10%, ** significant at 5% and *** significant at 1%. All regressions include district and survey fixed effects, an indicator of being within 20 km of a mine and farmer controls, which includes: household head's age, literacy, and an indicator of being born in the village; as well as an indicator of the household owning a farm plot. Columns 3 and 4 are estimated using 2SLS. The excluded instruments are: ln(area of land owned) and ln(number of adults equivalents in the household). Cumulative gold production is measured in hundreds of tonnes.

Figure 3: The effect of mining on agricultural productivity, by distance to a mine

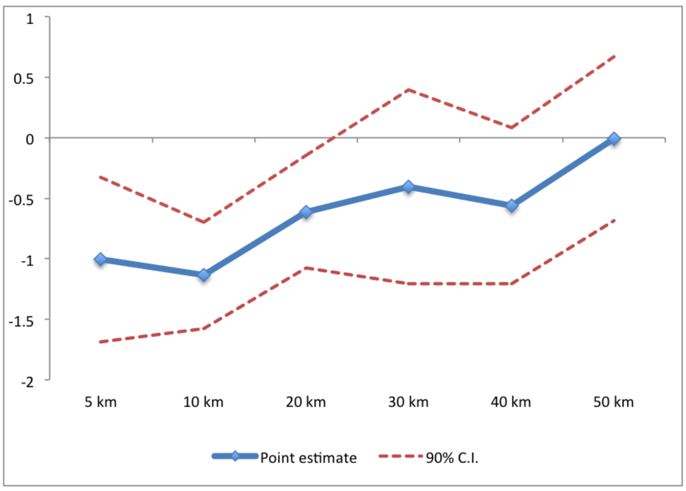

and Rosen (2012). This approach uses instrumental variables that *may be correlated to the error term.* Under weaker assumptions that the standard IV approach, this methodology allows us to identify parameters bounds instead of point estimates. We allow both instruments to be imperfect and run the IIV specification for different combinations of values of the parameters that measure the ratio of correlations of the instrument and the regressor with the error term, namely λ_{land} and λ_{labour}.[46] Figure 4 shows that the effect on residual productivity is negative in more than 95% of estimations. For all combinations where $\lambda_{land} < 0.5$ the corresponding estimate of the effect of pollution on agricultural output is negative. This suggests that the direction of the effect is insensitive to allowing the correlation between the land instrument and the error term to be up to half that of the correlation between actual land use and the error term.

4.1.1 Robustness checks

In Table 4 we check that our results are robust to alternative specifications.[47] Column 1 estimates a parsimonious model without farmer characteristics. Column 2 includes all controls and adds indicators of use of other inputs (such as chemical fertiliser, manure and improved seeds). Column 3 further expands this specification by adding an array of heterogeneous trends. We include the interaction of time trends with indicators of ecological zone, proximity to coast and to region capitals. This last specification addresses concerns that the measure of mining activity may be picking up other confounding trends.

Column 4 excludes observations within 5 km of a mine. This addresses concerns that the effects are driven by factors such as land grabbings and population displacement. Population displacements are usually confined to the mine operating sites, i.e. areas containing mineral deposits, processing units, and tailings. These areas comprise, at most, few kilometers around the mine site.

Column 5 performs a falsification test, where we estimate the baseline regression (2) including interactions between time trends and dummies of: (1) proximity to an active mine,

[46]Note that $(\lambda_{land}, \lambda_{labour}) = (0, 0)$ corresponds to the standard 2SLS estimate. For further details of the methodology see Nevo and Rosen (2012, section III.D).

[47]Results are also robust to the inclusion of mine fixed effects, exclusion of farmers in the vicinity of Obuasi mine, and use of a CES production function (see Tables A.8, and B.1 in the online Appendix). As discussed in Section 2, Obuasi mine's operations were of a sizable magnitude before the period of interest. The results of checks in Table 4 are similar using instruments for labour and land (see Table A.7 in online appendix).

Figure 4: Estimates of γ with multiple imperfect IVs

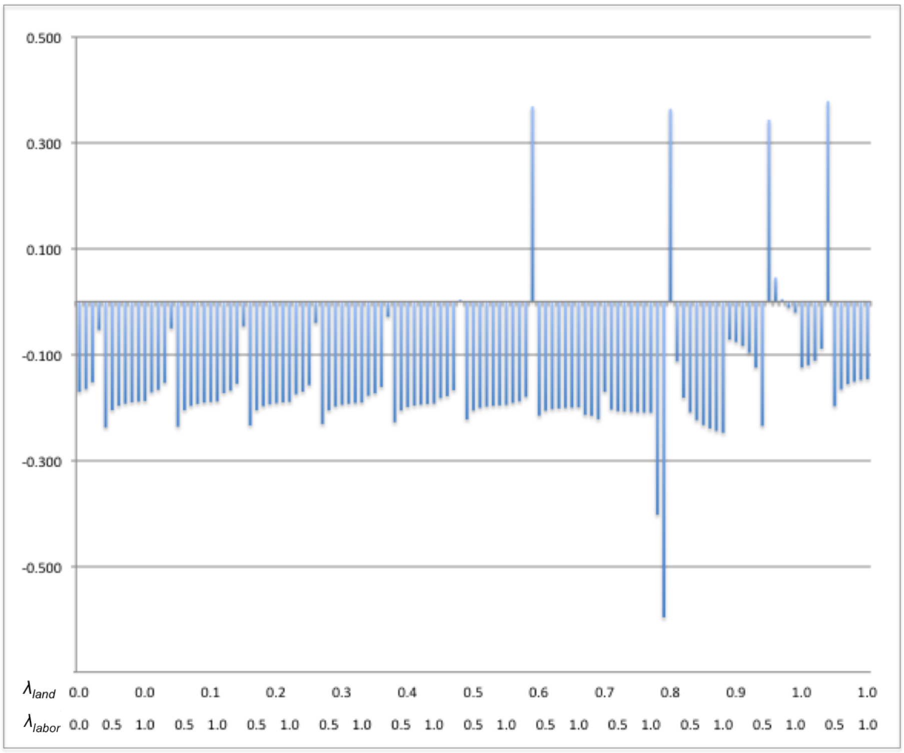

Note: Vertical axis displays estimates of γ for different values of λ_j, with $j = \{land, labour\}$. Values of λ_j in horizontal axis range from 0 to 1, with step increments of 0.1. $\lambda_j = \frac{\text{corr}(Z_j,\epsilon)}{\text{corr}(X_j,\epsilon)}$, where X is input use, Z is the instrumental variable and ϵ is the error term, measures how well the instrument satisfies the exogeneity assumption. $\lambda_j = 0$ corresponds to an exogenous, valid, instrument. The assumption that the instrument is less correlated to the error term that the endogenous variable implies that $\lambda_j < 1$.

and (2) proximity to a future mine, but not to an active one. Future mines include sites that started operations after 2005 or have not started production yet but are in the stage of advanced exploration or development. The results show that the negative relation between mining and agricultural productivity occurs only in the proximity of mines active during the period of analysis, but not in future mining areas.

4.2 Is this driven by pollution?

We interpret the previous findings as evidence that agricultural total factor productivity has decreased in the vicinity of mines. We argue that a plausible channel is through the presence of mining-related pollution. As discussed above, modern mines can pollute air with exhausts from heavy machinery and processing plants, and particulate matter from blasting. In low concentrations, these pollutants are dispersed and absorbed by the environment. In larger concentrations, however, they can deposit on the ground in the form of acid rain and thus have long-term cumulative effects. This is in addition to other industry specific pollutants, such as cyanide, heavy metals and acidic discharges, which may also have cumulative effects but are mostly dispersed through surface water.

To further explore this issue, ideally we would need measures of environmental pollutants at local level in order to examine whether mining areas are indeed more polluted. Unfortunately, this information is not available in the Ghanaian case.[48]

Instead we rely on satellite imagery to investigate whether there is evidence of pollution that may be attributed to mining activities.[49] The satellite imagery is obtained from the Ozone Monitoring Instrument (OMI) available at NASA, which provides daily measures of tropospheric air conditions since October 2004.[50] We focus on one particular air pollutant: nitrogen dioxide (NO_2). The negative effects of NO_2 can be both short-term, by directly damaging plant's tissues, or cumulative, through acid rain and the subsequent degradation of soils. The main source of

[48]An alternative way to assess exposure to pollution is to use information collected by Ghana's Environmental Protection Agency (EPA). This agency collects information of environmental pollutants in some mining areas, and produces environmental assessments. This information has, however, two main limitations. First, the information has been collected only since 2009; hence it may not accurately reflect the environmental conditions during the period of analysis (1997-2005). Second, there are not environmental assessments for all mines that were active before 2005, nor for non-mining areas that could be used as a control group. These issues limit their use in formal regression analysis.

[49]A similar approach of using satellite imagery to measure air pollutant is used by Foster et al. (2009) and Jayachandran (2009).

[50]For additional details, see `http://aura.gsfc.nasa.gov/instruments/omi.html`. Data are available at `http://mirador.gsfc.nasa.gov/cgi-bin/mirador/presentNavigation.pl?tree=project&project=OMI`.

Table 4: Robustness checks

	ln(real agricultural output)				
	(1)	(2)	(3)	(4)	(5)
Cumulative gold	-0.169*	-0.163*	-0.166*	-0.163*	
prod. within 20 km	(0.096)	(0.084)	(0.087)	(0.087)	
Within 20 km of active					-0.800***
mine × GLSS 5					(0.280)
Within 20 km of future					0.441
mine × GLSS 5					(0.435)
ln(land)	0.669***	0.599***	0.603***	0.637***	0.630***
	(0.039)	(0.039)	(0.039)	(0.039)	(0.038)
ln(labor)	0.220***	0.207***	0.206***	0.205***	0.212***
	(0.031)	(0.032)	(0.034)	(0.033)	(0.031)
Use fertiliser		0.444***	0.446***		
		(0.098)	(0.098)		
Use manure		0.548***	0.549***		
		(0.153)	(0.154)		
Use improved seeds		-0.108	-0.111		
		(0.092)	(0.090)		
Farmer controls	No	Yes	Yes	Yes	Yes
Heterogenous trends	No	No	Yes	No	No
Sample	All	All	All	Excl. within 5 km of mine	All
Observations	1,627	1,627	1,627	1,598	1,627
R-squared	0.422	0.464	0.465	0.448	0.454

Notes: Robust standard errors in parentheses. Standard errors are clustered at district level. * denotes significant at 10%, ** significant at 5% and *** significant at 1%. All regressions are estimated using OLS, and include district and survey fixed effects, and an indicator of being within 20 km of a mine. Column 2 replicates the baseline regression in Table 3 but includes indicators of other inputs, such as fertilisers, manure and improved seeds. Column 3 adds to the previous column the interaction of time trends with indicators of ecological zone, proximity to coast, and proximity to region capitals. Column 4 replicates the baseline regression but excludes farmers within 5 km of a mine. Column 5 performs a falsification test. *active mines* are mines that had some production in period 1988-2005, while *future mines* are mines that started operations after 2005 or have not started production yet, but are in the stage of advanced exploration or development.

NO_2 is the combustion of hydrocarbons such as biomass burning, smelters and combustion engines and is likely to occur near large urban centres, industrial sites and heavily mechanised operations, such as large-scale mines.

There are three important caveats relevant for the empirical analysis. First, the satellite data reflect air conditions not only at ground level where they can affect agriculture, but in the entire troposphere (from ground level up to 12 km).[51] Levels of tropospheric and ground level NO_2 are, however, highly correlated.[52] Thus, data from satellite imagery is informative of surface levels of NO_2. Second, the data is available only for 2005, the end of the period of analysis, therefore we can only exploit cross-sectional variation in air pollution. Finally, the measures of NO_2 are highly affected by atmospheric conditions, such as tropical thunderstorms, cloud coverage, and rain that are particularly important from November to March, and during the peak of the rainy season.[53] For that reason, we aggregate the daily data taking the average over the period April-June 2005, corresponding to the beginning of the rainy season and to the start of the main agricultural season.

To compare the relative levels of NO_2 in mining and non-mining areas, we match the satellite data to each enumeration area and estimate the following regression:[54]

$$NO2_v = \phi_1 X_v + \phi_2 W_v + \omega_v,$$

where $NO2_v$ is the average value of tropospheric NO_2 in enumeration area v during the period April-June 2005. X_v is an indicator of proximity to a mine and W_v is a vector of control variables.[55] Note that the unit of observation is the enumeration area and that, in contrast to the baseline results, this regression exploits cross-sectional variation only.

Column 1 in Table 5 presents the empirical results. We also replace the dummy X_v by a distance spline with breaks at 10, 20, 30 and 40 km and plot the resulting estimates in Figure 5, excluding farmers farther away.

[51] To obtain accurate measures at ground level, we would need to calibrate existing atmospheric models using air measures from ground-based stations. This information is unavailable.

[52] The correlation between these two measures is typically above 0.6. OMI tropospheric measures tend to underestimate ground levels of NO_2 by 15-30% (Celarier et al., 2008).

[53] In southern Ghana, the rainy season runs from early April to mid-November.

[54] The satellite data are binned to 13 km x 24 km grids. The value of NO_2 of each enumeration area corresponds to the value of NO_2 in the bin where the enumeration area lies.

[55] NO_2 is measured as 10^{15} molecules per cm^3. The average NO_2 is 8.1 while its standard deviation is 1.1.

The satellite evidence suggests that mining areas have a significantly greater concentration of NO_2. Moreover, the concentration of NO_2 decreases with distance to the mine in a similar fashion as the observed decline in total factor productivity. These latest findings point out to air pollution as a plausible explanation for the decline of agricultural productivity in mining areas.

Columns 2 further explores the relation between mining, air pollution and productivity. To do so, we estimate the relation between NO_2 and agricultural productivity using an indicator of proximity to a mine as an instrument for NO_2. Since we only have measures of NO_2 for 2005, we use the sample of farmers in the GLSS 5 and thus exploit only cross sectional variation. Consistent with mining-related pollution being a possible explanation, we find a significant negative correlation between NO_2 and agricultural productivity.[56]

So far, we have been using measures of the *stock* of pollutants, i.e. cumulative production. We use this measure due to the potential of many mining-related pollutants (such as air emissions and heavy metals) to have cumulative effects on the environment. Here, we check whether measures of the *flow* of pollutants would be better instead. As a measure of the flow of pollution, we use the annual production of the neighbouring mines in the surveys' reference years, i.e. 1997 and 2005. Columns 3 and 4 in Table 5 display the results. First, we add only the measure of flow of pollution. Then, we include both measures of stock and flow of pollution. The results suggest that the reduction in productivity is only affected by the variation in the measure of long-term exposure to pollution. This reflects the fact that cross-mine variation in production for the two relevant years is actually very small to drive our results.

Finally, we explore the importance of pollutants carried by surface water. To do so, we identify areas downstream of active mines and examine whether the negative effects of mining are stronger in these areas. Note that this is a crude way to assess exposure to pollution since some pollutants (such as heavy metals and dust) can be carried by both water and air, therefore areas upstream and downstream of mines can both be negatively affected.

We replicate the baseline regression including an interaction term between our measure of mining activity and a dummy variable *downstream* that is equal to one if the household is located downstream of an active mine. The results, displayed in column 5 of Table 5, suggest

[56]In the first stage, the relation between NO_2 and the excluded instrument (*within 20 km of mine*) is positive and significant at 5%.

Table 5: Mining and pollution

	Average NO_2	ln(real agricultural output)			
		Using mining as IV	Stock vs flow pollution		Upstream vs downstream
	(1)	(2)	(3)	(4)	(5)
Within 20 km of mine	0.325***				
	(0.111)				
Average NO_2		-1.554*			
		(0.837)			
Cumulative gold prod. within 20 km				-0.220**	-0.193**
				(0.093)	(0.094)
Annual gold prod. within 20 km			-0.057	1.644	
			(1.324)	(1.802)	
Cumul. gold prod. within 20km × downstream					-0.012
					(0.086)
Estimation	OLS	2SLS	OLS	OLS	OLS
Farmer controls	No	Yes	Yes	Yes	Yes
Controlling for inputs	No	Yes	Yes	Yes	Yes
Observations	399	914	1,627	1,627	1,627
R-squared	0.238	0.029	0.443	0.445	0.447

Notes: Robust standard errors in parentheses. * denotes significant at 10%, ** significant at 5% and *** significant at 1%. Standard errors in columns 2 to 5 are clustered at district level. Columns 1 and 2 use data for 2005 only. Column 1 uses the enumeration area as unit of observation and includes indicators of ecological zones, urban area, and region fixed effects as additional controls. Column 2 presents 2SLS estimates of the agricultural production function using only the sample of farmers in GLSS 5. It treats *Average* NO_2 as an endogenous variable and uses *within 20 km of mine* as the excluded instrument. This specification includes the additional controls: indicators of ecological zone, urban area, region fixed effects, as well as farmer characteristics and measures of input use as in the baseline regression (see notes of Table 3). Column 3 and 4 replicates the baseline OLS regression (column 2 in Table 3) adding *annual gold production within 20 km* as a proxy for flow pollutants. This variable measures the production of gold (in hundreds of MT) from nearby mines in years 1997 and 2005. Column 5 adds to the baseline OLS regression an interaction term of the measure of mining activity and *downstream*, a dummy equal to one if household is downstream of an active mine. This regression also includes *downstream* and its interaction with *within 20 km of mine*.

that there is no significant difference in the effect of mining between downstream or upstream areas. A conservative interpretation here is that pollution of surface waters may not be driving the main results, although this may be due to low statistical power.[57]

Figure 5: Average concentration of NO_2, by distance to a mine

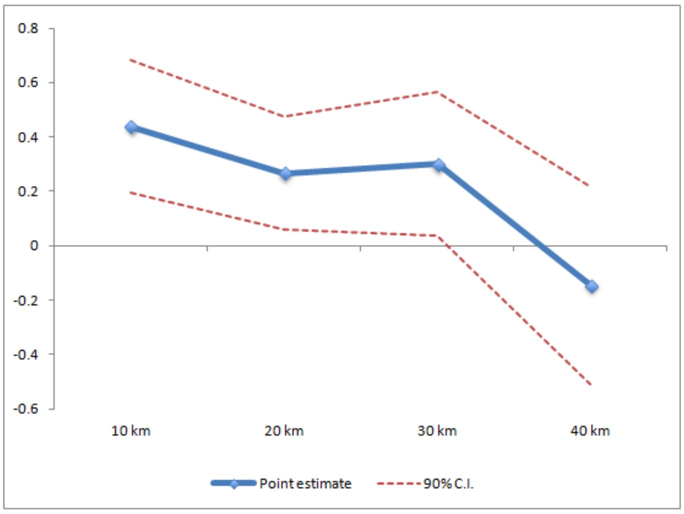

4.3 Alternative explanations

We interpret the previous results as evidence that pollution is a credible channel through which mining has affected agricultural productivity. In this section, we examine three plausible alternative explanations.

First, mining can directly appropriate some inputs, for instance by diverting water sources or the appropriation of farmland. A concern is that the drop in productivity simply reflects the relocation of farmers to less productive lands.[58] It is unlikely, however, that this factor fully accounts for the observed reduction in productivity as the effects we found are over a very large area (in excess of 1200 square km). Furthermore, column 4 in Table 4, shows that the results are robust to the exclusion of farmers within 5 km of a mine, the population most likely to suffer from displacement.[59]

[57]Additionally, there is no variation in productivity that can be explained by the direction of winds. Ghana has two main winds that come from opposite directions: the Harmattan, a dry and dusty wind, that blows from the Sahara, i.e. north east, and another wind, warm and moist, coming from the Atlantic ocean, i.e., south-west. Hence, air pollutants may be dispersed in all directions around a mine.

[58]These phenomena are documented in the Ghanaian case and are deemed to be a source of conflict and increased poverty in mining areas (Duncan et al., 2009; Botchway, 1998).

[59]We also examine the relation between mining and agricultural input prices, see Section D in online appendix. We find no evidence of an increase in agricultural wages or land prices in mining areas as compared to non-mining

Second, mines could be hiring local workers or fostering a local demand boom as documented in Aragon and Rud (2013) for a gold mine in northern Peru. This can attract workers away from agriculture towards mining or other non-tradable sectors. If these relocating workers are more productive, then the reduction in agricultural productivity would be just reflecting changes in the *composition* of agricultural workers. A similar phenomenon could occur in the presence of selective migration, for instance if more productive farmers migrate away from mining areas.

To assess this alternative explanation, Table 6 examines whether mining activity is associated with changes in several observable population characteristics. In columns 1 and 2, we look at the probability that a working-age individual, male or female, is employed, self-employed or engaged in domestic production. In column 3, we look at the probability that a worker is engaged in agriculture (either as a producer or labourer). In the presence of occupational change towards non-agricultural activities, we could expect a negative correlation. Columns 4 and 5 examine measures of agricultural workers' demographics and mobility, such as probability of being a prime age male (20-40 years), or being born in the same village where they reside. Finally, columns 6 and 7 explore measures of human capital of agricultural workers, such as literacy and having completed secondary school.[60] This result is informative, however, under the assumption that farming ability is positively correlated with educational attainment, which is a plausible assumption given that in our baseline regression the measure of literacy is associated with an increase in agricultural product and productivity. We find, however, no evidence of any change in these population characteristics.

Finally, an alternative story that could explain lower agricultural productivity is related to weak property rights. In Ghana, two phenomena are at play: customary and ill-defined land rights, and the right of the State to grant licenses for the use of land where mineral wealth is located (Botchway, 1998). Farmers near mining sites might fear expropriation and might choose to reduce agricultural investments, such as planting cocoa trees, as documented in Besley (1995). We first check whether there is a change in land ownership. Then, we examine whether there is any perceptible decrease in the share of cocoa or planting of new cocoa trees. Finally, we also explore changes in other agricultural practices, such as crop diversification and use of fertilisers,

areas.

[60] Levels of completion of primary school are high, i.e., around 86%, while literacy levels (47.8%) and secondary school completed (36.3%) show greater variation. Results hold when using data on completed primary school.

that could change as a way to mitigate the effect of pollution.[61]

Table 7 displays the results. We do not find a decrease in cocoa planting nor significant changes in land ownership or the use of fertilisers. If anything, there has been an increase in planting of cocoa trees. These results contradict the property rights explanation, and weaken the argument that the reduction in productivity is driven solely by changes in perceived risk of expropriation. Interestingly, we in fact find an increase in crop concentration. While far from conclusive, this finding is suggestive of actions taken by farmers to ameliorate the negative effect of pollution.

The findings discussed above, together with the observed increase of air pollution in the vicinity of mines, supports our finding that pollution is the most plausible channel for mining to affect agricultural productivity. These results, however, should not be interpreted as conclusive evidence that mining affects agriculture *only* through pollution, as other channels may also be important. For instance, a local mining boom may have changed the composition of workers in an unobservable dimension. Similarly, improvements in the outside options of agricultural workers (such as artisanal mining or urban services) may reduce their incentives to exert effort in the farm. However, due to data limitations we are unable to examine these explanations.

Table 6: Population characteristics

	Do any work (1)	Do any work (2)	Works in agriculture (3)	Male in prime age (4)	Born in village (5)	Literacy (6)	Completed secondary (7)
Cumulative gold prod. within 20 km	-0.001	-0.018	-0.032	-0.001	-0.006	-0.004	-0.013
	(0.006)	(0.017)	(0.042)	(0.018)	(0.024)	(0.021)	(0.016)
Sample	Males in working age	Females in working age	All workers	Agricultural workers	Agricultural workers	Agricultural workers	Agricultural workers
Observations	4,787	5,688	8,932	4,978	4,929	4,971	4,978
R-squared	0.453	0.319	0.359	0.029	0.127	0.044	0.134

Notes: Robust standard errors in parentheses. Standard errors are clustered at district level. * denotes significant at 10%, ** significant at 5% and *** significant at 1%. All regressions include district and survey fixed effects, an indicator of being in a mining area, and indicators of ecological zone and urban area. *Do any work* is an indicator equal to one if individual is employed, self-employed or participates in domestic production. Working age is between 15 to 65 years. *Works in agriculture* is an indicator equal to one if individuals works in agriculture as a laborer or producer. *Male in prime age* is an indicator equal to one if individual is male between 20 to 40 years old. *Born here* is an indicator equal to one if individual was born in the same village where she resides. All regressions are estimated using a linear probability model. Columns 1 to 3 include as additional controls: age, age^2, religion, place of birth, literacy status, and household size. Columns 6 and 7 examine the educational attainment of agricultural workers conditional on age and age^2.

[61] Farmers could ameliorate the effects of soil degradation by increasing the use of fertilisers. Similarly, if crop sensitivity to pollution is heterogeneous, farmers may reduce the impact on their income by changing the composition of crops farmed.

Table 7: Agricultural investment and practices

	Owns farm (1)	New cocoa plants (2)	Share of cocoa (3)	Crop concentration (4)	Use fertiliser (5)	Use manure (6)
Cumulative gold prod. within 20 km	-0.010	0.066*	0.021	0.043**	-0.005	-0.015
	(0.029)	(0.039)	(0.036)	(0.017)	(0.045)	(0.026)
Observations	1,627	1,627	1,627	1,627	1,627	1,627
R-squared	0.225	0.159	0.446	0.118	0.140	0.102

Notes: Robust standard errors in parentheses. Standard errors are clustered at district level. * denotes significant at 10%, ** significant at 5% and *** significant at 1%. All regressions include district and survey fixed effects, an indicator of being in a mining area, and indicators of ecological zone and urban area. Columns 2 to 6 also include farmer's controls as the agricultural production function in Table 3. All regressions are estimated using linear OLS. *Owns farm* is equal to 1 if farmer owns any plot. *New cocoa plants* equals one if the farmer has planted new cocoa trees in the previous 12 months. *Share of cocoa* is the share of cocoa revenue in the value of total agricultural output. *Crop concentration* is the Herfindahl concentration index of crops' value. Outcomes in columns 5 and 6 are indicators equal to one if farmer uses chemical fertilisers or manure, respectively.

4.4 Exploring the mechanisms

An important question concerns the exact mechanism by which mining-related pollution affects total factor productivity. As discussed in Section 2.2, we consider three possible mechanisms. First, pollution could directly affect crop yields and health. Second, pollution could deteriorate the quality of key inputs, such as soil. Third, through its effects on human health, pollution could affect labour productivity.

To formally discuss these factors, we consider the following augmented Cobb-Douglas production function:

$$Y = q_T(q_M M)^{\alpha}(q_L L)^{\beta} \tag{3}$$

where Y is agricultural output, M and L are the observable quantities of land and labour. q_L and q_M are input-specific quantity shifters, which respectively capture factors such as labour productivity and quality of soil, while q_T captures all other unobserved factors, including crop health and yields. Our previous discussion suggests that pollution could potentially affect any of these factors.

In this setup, total factor productivity is captured by $A = q_T q_M^{\alpha} q_L^{\beta}$. This is the object that we can observe, as a residual, when we estimate the agricultural production function. Our empirical analysis shows that mining-related pollution reduces A but with the data at hand we cannot identify its effect on each component as this would require data on quality of soil, crops' health and labour productivity.

Instead, we use an indirect approach to show that the negative effect of mining on total factor productivity cannot be entirely driven by reduction in labour productivity, q_L. We do this in three ways: First, we note that under the assumption that all the reduction in A is driven solely by changes in labour productivity implies that $\Delta \ln A = \beta \Delta \ln q_L$. In Table 3 column 4, we estimate $\Delta \ln A = -0.565$ and $\beta = 0.346$. This suggests a reduction in q_L of almost 80%. However, this figure is inconsistent with previous estimates of the relation between air pollution and labour productivity. For instance, Graff Zivin and Neidell (2012) find that one standard deviation in ozone levels decreases labour productivity by roughly 5.5% using U.S. data.

Second, we examine worker health indicators. We use self-reported data on the incidence and duration of illness.[62] We then examine the relation between these measures of health and our measure of mining. We focus on working age individuals (aged 15 to 65) and split the sample between urban and rural populations. Table 8 displays the results. In all cases, we find no evidence of an increase in the likelihood of being ill nor on the duration of illness. This is contrary of what we could expect if the sole channel was through human health.

Table 8: Mining and self-reported illness

	Ill in previous 2 weeks			ln(number of days ill)		
	(1)	(2)	(3)	(4)	(5)	(6)
Cumulative gold prod. within 20 km.	-0.015	0.013	-0.022	0.019	-0.182***	0.038
	(0.022)	(0.046)	(0.026)	(0.032)	(0.034)	(0.034)
Sample	All	Urban	Rural	All	Urban	Rural
Observations	11,713	4,498	7,215	2,842	1,041	1,801
R-squared	0.055	0.066	0.071	0.062	0.089	0.081

Notes: Robust standard errors in parentheses. Standard errors are clustered at district level. * denotes significant at 10%, ** significant at 5% and *** significant at 1%. All regresssions are estimated using OLS and include district and survey fixed effects, an indicator of being within 20 km of a mine, and individual controls such as: age, age^2, gender, an indicator of rural area and ecological zone. *Ill in previous 2 weeks* is a dummy variable equal to 1 if individual reports being ill during the last 2 weeks, which does not include accidents. Column 4 to 6 include only the subset of individuals who reported being ill.

Finally, we examine the effect of mining on urban workers, not directly linked to the agricultural sector. This group include employed and self-employed workers. We focus on two available outcomes: number of hours worked and employment income. Under reasonable assumptions,

[62]The survey questions are: In the last two weeks, have you been ill? If yes, how many days have you been ill?.

if the effect was transmitted entirely through reduction in labour productivity, we should also observe a decrease in these labour outcomes.[63] The results are displayed in Table 9. Column 1 and 3 use the sample of all urban workers, including agricultural workers, while columns 2 and 4 include only non-agricultural workers. Note that all regressions exclude mining workers, who can be directly affected by mining operations. In all cases, there is no significant change in number of hours nor on employment income.

Taken together, this evidence does not rule the possibility that the effects reflect, in part, reduction in labour productivity. However, they suggest that it is unlikely that this mechanism fully accounts for the observed phenomena.

Table 9: Mining and labor outcomes of urban workers

	ln(hours work)		ln(real employment income)	
	(1)	(2)	(3)	(4)
Cumulative gold prod. within 20 km.	-0.062	-0.064	0.222	0.139
	(0.042)	(0.064)	(0.260)	(0.250)
Sample	All urban workers	Urban non-agric. workers	All urban workers	Urban non-agric. workers
Observations	2,580	2,062	1,936	1,564
R-squared	0.152	0.090	0.389	0.319

Notes: Robust standard errors in parentheses. Standard errors are clustered at district level. * denotes significant at 10%, ** significant at 5% and *** significant at 1%. All regressions include district and survey fixed effects, an indicator of being in a mining area, and indicators of ecological zone and urban area. Columns 1 and 2 include as additional controls: age, age^2, religion, place of birth, literacy status, and household size. Columns 3 and 4 add as additional control the log of number of hours worked. All regressions exclude mining workers. Columns 2 and 4 also exclude agricultural workers.

5 Effects on poverty

The standard consumer-producer household framework presented above links a household's utility function, which depends on consumption levels, to income from agricultural production. As a consequence, we expect that our previous results indicating a sizable reduction in agricultural productivity and output induce a subsequent effect on local living standards, such as measures of poverty. There are reasons to believe that this channel can be averted. Mining companies or

[63] These assumptions are: (1) labour demand for urban workers depend of their productivity, (2) mining did not increase labour demand in urban areas, and (3) mining did not affect urban labour supply. The first assumption is more reasonable given the existence of urban labour markets. The last two assumptions are likely to be met given the limited economic interactions between gold mines and local economies in the Ghanaian context.

the government could, for example, promote local development projects, employ local workers, compensate local residents, or transfer part of the mining surplus. These policies are often implemented by the industry to mitigate potential negative side-effects of mining, and may offset the decline in productivity.

To examine this issue, we use data from the GLSS on poverty and estimate a difference-in-difference regression of household poverty on our measure of mining activity, S_{vt}.[64] The results are displayed in Table 10. Column 1 shows results for all households using our preferred specification with cumulative gold production as a measure of mining activity. As a reference, column 2 uses as proxy of S_{vt} the interaction between a dummy of proximity to a mine and a time trend to obtain the average effect of mining on poverty. Columns 3 and 6 split the rural sample between urban and rural households, respectively. Column 4 looks at rural households that are engaged in household production (and thus were included in the estimation of the agricultural production function), while column 5 looks at rural households that did not report any agricultural production.[65] We also check the robustness of the results to using a continuous measure of real household expenditure (see table E.1 in the online appendix).

The picture that emerges is that there is a positive and significant relation between mining activity and poverty. The magnitude of the effect is sizable: the increase in gold production between 1997 and 2005 is associated with an increase of almost 18 percentage points in poverty headcount. The effect is concentrated among rural inhabitants, regardless of whether the households are agricultural producers or not. Non-producers could be affected either directly, by the reduction in agricultural wages associated to lower total factor productivity, or indirectly, if they sell good or services to local farmers.[66]

The reduction in indicators of economic well-being is consistent with the decline in agricultural productivity in areas where farming activities are the main source of livelihood. Table E.2 in the online appendix shows two additional results among children that are also consistent with levels of poverty induced by pollution: malnutrition and acute respiratory diseases have both increased in mining areas.

These results, however, should not be interpreted as evidence that mining affects local

[64]See section E in online appendix for further estimation details and results.

[65]Households whose members are engaged in farming as wage labourers are around 65% of the sample.

[66]Aragon and Rud (2013) discuss the conditions under which these effects would be present and show evidence of how households were affected in the area of influence of a gold mine in Peru.

economic conditions *only* through its effect on agriculture. Mining could have created a local demand shock, affected provision of public goods, or changed the scope of other re-distributive policies. Similarly, mining could generate other local negative effects, such as an increase in rent-seeking, conflict, or political corruption (Caselli and Michaels, 2013). Despite these limitations, these results are indicative of the *net* effect of mining on local living conditions. Compensating policies and positive spillovers from mines, if any, have been insufficient to offset the negative effect on agricultural income.

Table 10: Mining and poverty

	Poverty					
			Rural			Urban
	All households		All	Farmers	Non-farmers	
	(1)	(2)	(3)	(4)	(5)	(6)
Cumul. gold prod. within 20 km.	0.059***		0.071***	0.056**	0.084**	0.054
	(0.015)		(0.019)	(0.021)	(0.032)	(0.036)
Within 20 km of mine × GLSS 5		0.186***				
		(0.055)				
Observations	5,527	5,527	3,393	2,540	853	2,134
R-squared	0.212	0.216	0.227	0.237	0.224	0.199

Notes: Robust standard errors in parentheses. Standard errors are clustered at district level. * denotes significant at 10%, ** significant at 5% and *** significant at 1%. All regressions are estimated using ordinary least squares, and include district and survey fixed effects as well as household controls, such as: age, age^2, religion, place of birth and literacy status of household head, household size, and an indicator of urban areas. All columns include an indicator of being within 20 km of a mine.

6 Concluding remarks

This paper examines an important externality that polluting industries may impose on rural areas, namely, a reduction in agricultural productivity. We find robust evidence that agricultural productivity has decreased in mining areas relative to areas in the same region but located at a greater distance from mining activities. The reduction is economically significant: approximately a 40% decline in total factor productivity between 1997 and 2005. We also document an increase in rural poverty associated to the decline in agricultural productivity. The magnitude of these effects is, however, specific to the Ghanaian case we study and should not be extrapolated to other contexts.

These findings have an important implication for environmental and industrial policies. In particular, they suggest that environmental assessments should consider the possible impact of polluting industries on agricultural productivity and farmers' income.

These potential costs are usually neglected in the academic and policy debate, which usually focuses on the benefits extractive industries could bring in the form of jobs, taxes or foreign currency. These benefits are weighted against environmental costs such as loss of biodiversity, or human health risks. However, local living standards may be also directly affected by the reduction in agricultural productivity. In fertile rural environments, these costs may offset the benefits from extractive industries, and hinder the ability to compensate affected populations. In turn, this may have substantial re-distributive effects.

A simple back of the envelope using the Ghanaian case illustrates this argument. In 2005, mining-related revenues amounted to US$ 75 million, which represent around 2-3% of total government revenue. Most of this revenue (around 80%) was channeled to the central government.[67] In contrast, the average annual loss by farming households in mining areas, according to our main results, is in the order of US$ 97 million.[68] These approximate numbers show that the amount of tax receipts might not be enough to compensate those farmers negatively affected by mining and that this situation is even worsened by the fact that only a small proportion of the tax receipts go back to affected localities.

A main limitation of this paper is that we cannot clearly assess the relative importance of several plausible mechanisms through which pollution could affect productivity, such as the direct effects of pollution on labour productivity, quality of soil, and crop health. Similarly, we cannot examine in detail changes in farmers' decisions to ameliorate the effect of pollution. While beyond the scope of this paper, examination of these issues warrant further research.

[67] Local authorities (such as District Assemblies, Stools and Traditional Authorities) receive only 9% of mining royalties.

[68] This number is obtained by multiplying the number of producing households in mining areas, around 210,000, to the average reduction in households' per capita annual consumption, i.e., US$ 460.

Acknowledgements

This project was funded by the International Growth Centre, grant RA-2010-12-005. We thank Douglas Gollin, Mushfiq Mubarak, Matt Neidell, Owen Ozier, Krishna Pendakur, Francis Teal, Chris Udry, Tanya Wilson and seminar participants at BREAD/World Bank, Calgary, CSAE, EAERE, EEA-ESEM, IGC Ghana, IGC Growth Week, Novafrica, Grantham Research Institute-LSE, Oxcarre, Gothenburg, NEUDC, Tinbergen Institute, Univ. of Washington, Univ. of Oslo, Paris 1 Pantheon - Sorbonne, Paris School of Economics, Namur, UC Santa Barbara, UC Riverside, and Univ. of British Columbia for useful comments and suggestions. Godwin Awuah and Sarah Khan provided excellent research assistance. We also thank Phylomena Nyarko at Ghana Statistical Service and Henry Telly at IGC Ghana for invaluable help with this project.

References

Ackerberg, Daniel, Kevin Caves, and Garth Frazer, "Structural identification of production functions," MPRA Paper 38349, University Library of Munich December 2006.

Aragon, Fernando M. and Juan Pablo Rud, "Natural resources and local communities: evidence from a Peruvian gold mine," *American Economic Journal: Economic Policy*, May 2013, *5* (2).

Armah, F. A., S. Obiri, D. O. Yawson, A. N. M. Pappoe, and A. Bismark, "Mining and Heavy Metal Pollution: Assessment of Aquatic Environments in Tarkwa (Ghana) using Multivariate Statistical Analysis," *Journal of Environmental Statistics*, 2010, *1* (4).

Aryeetey, Ernest, Osei Bafour, and Daniel K. Twerefou, "Impact of Mining Sector Reforms on Output, Employment and Incomes in Ghana, 1980-2002," Technical Publication 75, Institute of Statistical, Social and Economic Research (ISSER), November 2007.

Bardhan, Pranab and Christopher Udry, "Development Microeconomics," *Oxford University Press*, 1999.

Benjamin, Dwayne, "Household Composition, Labor Markets, and Labor Demand: Testing for Separation in Agricultural Household Models," *Econometrica*, 03 1992, *60* (2), 287–322.

Besley, Timothy, "Property Rights and Investment Incentives: Theory and Evidence from Ghana," *The Journal of Political Economy*, October 1995, *103* (5), 903–937.

Botchway, F. Nii, "Land Ownership and Responsibility for the Mining Environment in Ghana," *Natural Resources Journal*, Fall 1998, *38*, 509–536.

Caselli, Francesco and Guy Michaels, "Do oil windfalls improve living standards? Evidence from Brazil," *American Economic Journal: Applied Economics*, 2013, *5* (1), 208–238.

Celarier, E. A., E. J. Brinksma, J. F. Gleason, J. P. Veefkind, A. Cede, J. R. Herman, D. Ionov, F. Goutail, J.-P. Pommereau, J.-C. Lambert, M. van Roozendael, G. Pinardi, F. Wittrock, A. Schönhardt, A. Richter, O. W. Ibrahim, T. Wagner, B. Bojkov, G. Mount, E. Spinei, C. M. Chen, T. J. Pongetti, S. P. Sander, E. J. Bucsela, M. O. Wenig, D. P. J. Swart, H. Volten, M. Kroon, and P. F. Levelt, "Validation of Ozone Monitoring Instrument nitrogen dioxide columns," *Journal of Geophysical Research (Atmospheres)*, May 2008, *113* (D15), 15.

Currie, Janet, Eric A. Hanushek, E. Megan Kahn, Matthew Neidell, and Steven G. Rivkin, "Does Pollution Increase School Absences?," *The Review of Economics and Statistics*, November 2009, *91* (4), 682–694.

——, **Joshua S Graff Zivin, Jamie Mullins, and Matthew J Neidell**, "What Do We Know About Short and Long Term Effects of Early Life Exposure to Pollution?," Technical Report, National Bureau of Economic Research 2013.

Dudka, Stanislaw and Domy C Adriano, "Environmental impacts of metal ore mining and processing: a review," *Journal of Environmental Quality*, 1997, *26* (3), 590–602.

Duncan, Edward E., Jerry S. Kuma, and Seth Frimpong, "Open Pit Mining and Land Use Changes: An Example from Bogosu-Prestea Area, South West Ghana," *The Electronic Journal of Information Systems in Developing Countries*, 2009, *36*, 1–10.

Emberson, L.D., M.R. Ashmore, F. Murray, J.C.I. Kuylenstierna, K.E. Percy, T. Izuta, Y. Zheng, H. Shimizu, B.H. Sheu, C.P. Liu, M. Agrawal, A. Wahid, N.M. Abdel-Latif, M. van Tienhoven, L.I. de Bauer, and M. Domingos, "Impacts of Air Pollutants on Vegetation in Developing Countries," *Water, Air, & Soil Pollution*, 2001, *130*, 107–118.

Foster, Andrew, Emilio Gutierrez, and Naresh Kumar, "Voluntary Compliance, Pollution Levels, and Infant Mortality in Mexico," *The American Economic Review*, 2009, *99* (2), 191–197.

Graff Zivin, Joshua S. and Matthew J. Neidell, "The Impact of Pollution on Worker Productivity," *American Economic Review*, December 2012, *102* (7), 3652–3673.

—— **and** ——, "Environment, Health and Human Capital," 2013. unpublished.

Hanna, Rema and Paulina Oliva, "The Effect of Pollution on Labor Supply: Evidence from a Natural Experiment in Mexico City," NBER Working Paper No. 17302 August 2011.

Heck, Walter W, OC Taylor, Richard Adams, Gail Bingham, Joseph Miller, Eric Preston, and Leonard Weinstein, "Assessment of crop loss from ozone," *Journal of the Air Pollution Control Association*, 1982, *32* (4), 353–361.

Jayachandran, Seema, "Air Quality and Early-Life Mortality: Evidence from Indonesia's Wildfires," *Journal of Human Resources*, 2009, *44* (4), 916–954.

Kotsadam, Andreas and Anja Tolonen, "Mineral Mining and Female Employment," *Oxcarre Research Papers 114, Oxford University*, 2013.

Maggs, R., A. Wahid, S. R. A. Shamsi, and M. R. Ashmore, "Effects of ambient air pollution on wheat and rice yield in Pakistan," *Water, Air, Soil Pollution*, 1995, *85*, 1311–1316.

Marshall, Fiona, Mike Ashmore, and Fiona Hinchcliffe, "A hidden threat to food production: Air pollution and agriculture in the Developing World," Gatekeeper Series SA73, International Institute for Environment and Development (IIED) 1997.

Menz, Fredric C. and Hans M. Seip, "Acid rain in Europe and the United States: An update," *Environmental Science & Policy*, 2004, *7* (4), 253–265.

Miller, Joseph E, "Effects on photosynthesis, carbon allocation, and plant growth associated with air pollutant stress," in "Assessment of crop loss from air pollutants," Springer, 1988, pp. 287–314.

Nevo, Aviv and Adam M Rosen, "Identification with imperfect instruments," *Review of Economics and Statistics*, 2012, *94* (3), 659–671.

Salomons, W, "Environmental impact of metals derived from mining activities: Processes, predictions, prevention," *Journal of Geochemical Exploration*, 1995, *52* (1), 5–23.

Tetteh, Samuel, Adem A. Golow, David K. Essumang, and Ruphino Zugle, "Levels of Mercury, Cadmium, and Zinc in the Topsoil of Some Selected Towns in the Wassa West District of the Western Region of Ghana," *Soil and Sediment Contamination: An International Journal*, 2010, *19* (6), 635–643.

U.S. Environmental Protection Agency, "Acid rain," Website `http://www.epa.gov/acidrain/effects/forests.html`, accessed on 11 March 2014. 2012.