Women directors and cost efficiency*

Anastasia Litina[†] Luca J. Uberti [‡] Skerdilajda Zanaj[§]

March 2024

Abstract

Previous studies have shown that women directors exert a higher monitoring and audit effort than their male peers. A possible consequence of tougher monitoring by gender-diverse boards is a more efficient use of the company's resources. In this paper, we test this hypothesis by using a stochastic frontier methodology. First, we estimate cost efficiency in a representative sample of gold mines around the world. Next, we identify the impact of female directors on cost efficiency by exploiting a plausibly exogenous component of variation in female boardroom representation driven by peer effects. We find that an increase in female representation on the parent company's board translates into substantial efficiency gains across the mining operations controlled by the parent company. A one standard-deviation increase in the share of female directors increases cost efficiency by 0.12 standard deviations of our main efficiency index. This finding is robust to using alternative instruments for female representation, different stochastic-frontier methodologies, and alternative specifications that control for various correlates of gender boardroom diversity. Yet, we find that the efficiency gains induced by female directors do not necessarily lead to improvements in overall performance as measured by accounting profitability, in line with previous findings in the literature. Our results provide additional evidence of a distinctly female style in corporate leadership.

Keywords: Gender; Boards of directors; Cost efficiency; Stochastic Frontier Analysis; Mining

JEL codes: D22, D24, G39, M14

---

*For comments and suggestions on our paper, we are grateful to Arnaud Bourgain, Luisito Bertinelli, Piergiovanna Natale, Laura Pagani, Daniella Puzzello and Eva Sierminska. We would also like to thank participants at the 2nd Gender and Economics Workshop at the University of Luxembourg, June 2023. The usual disclaimer applies.

[†] University of Macedonia, Thessaloniki, Greece (corresponding author). Email: `alitina.uoi@gmail.com`

[‡] University of Milan-Bicocca, Milan, Italy. Email: `lucajacopo.uberti@unimib.it`

[§] University of Luxembourg, Luxembourg. Email: `skerdilajda.zanaj@uni.lu`

1 Introduction

Recent decades have seen a dramatic increase in female educational attainment and labour force participation across advanced economies. Yet, women's representation in economic leadership positions has lagged behind. To break this invisible but powerful 'glass ceiling', many countries have taken steps to promote gender-balanced representation on corporate boards. Besides obvious considerations of equality and fairness, the economic rationale of quota (and other related) policies is that gender-diverse boards improve corporate performance and increase firm value. The diversity argument claims that heterogeneous directors bring a variety of perspectives, skills and experiences to the board. The larger information endowment that results can promote superior outcomes − including more robust deliberation, better decision-making, and a greater ability to identify and implement innovative strategies (Harrison and Klein, 2007).

In this paper, we use firm- and plant-level data to examine how gender diversity on corporate boards affects a particular dimension of corporate performance that has received little or no attention in the literature – namely, the cost-efficiency of the firm's operations. Most previous studies have focused on how gender diversity affects the firm's *overall* performance (see Johnson et al. [2013]; Post and Byron, [2015] for systematic reviews). The results are generally mixed. While some papers find a positive effect of female directors on profitability and market value (Green and Homroy 2018), a relative majority reports either negative (Ahern and Dittmar 2012; Matsa and Miller 2013) or null effects (Gregory-Smith et al. 2014; Ferrari et al. 2022).

In response to these mixed findings, recent studies have turned to investigating the effects of boardroom gender diversity on more specific performance and corporate governance outcomes. While the inclusion of women in top decision-making bodies may not necessarily increase the bottom line, it may still affect the style of corporate leadership, bringing about changes in the way firms are governed and operate. In particular, previous studies have shown that female executives are more risk averse (Sah et al. 2022), undertake fewer acquisitions (Levi et al. 2014), and allocate more effort to monitoring the firm's management (Adams and Ferreira 2009; Nekhili et al. 2020). A possible consequence of a leadership style focused on risk containment and tough monitoring is an improvement in the efficiency of the firm's *existing* operations. Measuring efficiency, however, is not straightforward; partly for this reason, the impact of female directors on this particular performance outcome has received insufficient attention in the literature.

This paper uses a newly released global dataset on mining companies and a stochastic frontier methodology to fill this gap. To the best of our knowledge, we are the first to investigate the impact of female directors on efficiency outcomes using highly granular cost and output data at the plant level. In the analysis, we use a global representative sample of 136 mining companies operating 294 gold mines worldwide. The sample comes from proprietary data assembled by a reputed data-analytics firm (Glacier Rig Ltd.). The information available allows us to estimate reliably the cost-efficiency of individual gold mines. Cost-efficiency is defined as the ability of a productive unit (for example, a mine) to achieve the minimum feasible level of expenditure that is necessary to produce a bundle of outputs, given the market price of inputs and the technology in use. To separate the impact of different mining technologies from the efficiency with which a given technology is operated, we focus on a sample of mines which can be assumed to be using the same production process − namely, gold mines.

In the analysis, we exploit several interesting features of our data. First, we are able to match information at the company- and board-level (on female representation) with detailed accounting data (on output and production costs) at the mine level. Second, we can distinguish between three different cost categories, and therefore estimate three complementary dimensions of cost efficiency across individual mines. Third, the time coverage of the data (2012-2020) makes it possible to exploit variation in board composition within companies over time, and therefore control for company- and board-level unobservables.

The analysis proceeds in two steps. In the first step, we use the mine-level information to estimate well-specified cost functions (based on our three cost metrics) in a stochastic frontier framework (Battese and Coelli 1992). Any observed (upward) deviation from the lowest feasible costs of production implied by the mine's estimated cost function can be attributed to mine-specific inefficiencies. These deviations can be obtained in the form of residuals in post-estimation, leading to reliable indicators of cost inefficiency at the mine level.[1] In the second step, we relate our indicators of cost inefficiency to the share of women directors sitting on the board of the mine's parent company. OLS regressions indicate that an increase in female participation on the parent company's board is associated with modest but significant efficiency gains across the mining operations controlled by the parent company.

There are reasons to doubt whether these simple OLS relationships reflect the causal effect of female directors on the efficiency of the company's operations. *First*, unobserved characteristics of the company or its board (for example, corporate culture) may affect both the company's efficiency performance and the gender composition of the board. For instance, being sensitive to ethical and social-responsibility considerations may cause firms to be both more female-friendly *and* also to achieve more efficient operations. We address this concern throughout the analysis by using company fixed effects (FE), which provide estimates based on within-company variation in female representation over time.

Second, the observed relationships may be driven by selection effects. Less efficient companies may be less likely to select female directors − for instance, if an urgent need to improve efficiency crowds out considerations of fairness and equality. Conversely, prospective female directors may tend to self-select into more efficient companies, especially if (observed) efficiency is a signal for female-friendliness. To mitigate this concern in an OLS framework, we include lagged measures of company-level performance in our fixed-effects regressions. Information on ROA and asset turnover, in particular, is easily observable on annual reports and may be used by prospective female directors to select target companies.

Third, and most problematically, boards may be chosen optimally to maximize efficiency. Here, the direction of the bias is not clear a priori. If women directors are on average less experienced or less able than men, an

[1] We do not have the data to estimate production functions and obtain measures of 'technical efficiency'.

exogenous increase in female representation would impose a constraint on efficiency maximisation. Thus, the 'true' effects of female directors on cost efficiency should be smaller than implied by the OLS estimates (Ahern and Dittmar 2012). If on the other hand, women are subject to statistical or taste-based discrimination in the executive labor market (with female candidates having to display substantially higher levels of ability than otherwise similar male candidates to secure the same directorship), then an exogenous increase in female participation would *relieve* a constraint on efficiency-maximising recruitment (Comi et al. 2020: 6). In this case, the exogenous entry of women would bring more able directors to the board, and the 'true' effects of female directors on cost efficiency should be larger than implied by OLS.

To address this third concern (while also mitigating the second one), we turn to an identification strategy based on instrumental variables (IV). We suggest that the choice to appoint female directors may be subject to peer effects (see e.g., Bulow et al. [1985], for a theoretical grounding, and Bustamante and Fresard [2021], for empirical evidence). In particular, firms may be under pressure to hire more women directors if their direct competitors embark on efforts to promote gender equality in the boardroom. Peer pressure is likely to be strongest amongst peer firms headquartered in the same city, where formal and informal contact between mining executives is frequent, and horizontal cultural transmission likely to be operative. Thus, we use the average share of female directors amongst the sampled firms headquartered in the same city as an instrument for a firm's own share of female directors. We argue that the resulting IV estimates are more likely to reflect the impact of an exogenous change in female representation than the corresponding OLS estimates. In line with labour-market discrimination as a potential confounding mechanism, we find that the IV estimates indicate *larger* (and now substantial) efficiency gains from female boardroom representation than the corresponding OLS estimates.

Although our data do not allow us to investigate in detail the mechanism underlying this relationship, we show that our results are unlikely to be driven by an increase in other efficiency-enhancing characteristics typically associated with incoming female directors − such as independence from management and younger age. Rather, we suggest based on previous findings in the literature that incoming female directors have systematically different preferences regarding monitoring compared to their male peers (Adams and Ferreira 2009; Nekhili et al. 2020).[2] The different leadership culture that women directors bring to the board is likely to be consequential from the point of view of efficiency, and thus explain our findings.

In additional results, we show that the cost-efficiency gains induced by gender-diverse boards, albeit important economically, do not make the companies that control the mines in our sample any more profitable. In sum, our results are consistent with women directors bringing to the board a distinct style of leadership focused on tough monitoring and careful auditing of expenditures, without necessarily improving the overall performance of the firm. In this regard, the evidence we present is consistent with previous studies finding mixed effects of female directors on accounting performance and overall firm value (Ahern and Dittmar 2012; Gregory-Smith et al. 2014).

[2]Previous studies (Abraham 2023) found that women may harbor distinct perspectives regarding performance, extending even to their self-perceptions, a factor that could be linked to the level of monitoring effort they exert.

The paper is structured as follows. Section 2 motivates the analysis by reviewing the related literature, while highlighting our contribution to it. Section 3 discusses the data. Section 4 presents the empirical methodology and the empirical results. Section 5 concludes.

2 Related Literature and Contribution

Boards of directors serve two main roles: advising managers on important decisions; and monitoring them to ensure that they act in the best interest of shareholders (Adams and Ferreira 2009). Most of the empirical literature on female directors has focused on gender differences in executive positions, examining whether female participation improves corporate performance and firm value by strengthening the board's advisory and monitoring functions (Wolfgang et al. 2023).

Some studies exploit the exogenous introduction of board gender quotas for identification. Ahern and Dittmar (2012) find that the Norwegian quota law (2003) led to younger and less experienced boards, causing a large decline in Tobin's Q for affected firms. Exploiting the same quota law, Matsa and Miller (2013) argue that the firms that experienced an exogenous increase in female board participation are less willing to undertake workforce reductions and have higher total labor costs than control-group firms, with negative consequences for performance as measured by the ratio of operating profits to assets. The authors argue that these findings cannot be attributed to a decline in director ability, but to a distinctly 'female style in corporate leadership', in line with the mantra that 'women take care, while men take charge'.

The findings of Ahern and Dittmar (2012) and Matsa and Miller (2013) confirmed previous evidence based on panel regressions and US data (Adams and Ferreira 2009), but were qualified by subsequent work. Using a sample of French firms during 2001-2010, Bennouri et al. (2018) link female directors to improvements in corporate performance (as measured by ROA and ROE), but deteriorations in market value (as measured by Tobin's Q). Using the gender of CEOs' children as a source of exogenous variation in female director appointments, Green and Homroy (2018) report positive effects of board gender diversity on both accounting- and market-based measures of performance for a sample of European firms. Comi et al. (2020) find mixed results of board gender quotas across European countries: an overall negative effect in France and Spain, and a positive effect on labour productivity and TFP in Italy. Carbonero et al. (2021) show that the Italian quota law improved the export capabilities of affected firm. Ferrari et al. (2022), however, confirmed Comi et al.'s (2020) finding of a null effect of the Italian quota law on overall corporate performance as measured by ROA and Tobin's Q.

Our paper provides a complementary angle to this literature by studying a global sample of highly homogenous firms active in the same sub-sector (gold mining), rather than a heterogenous sample of firms from the same country. We focus on the mining industry, a notoriously male-dominated sector where executives and directors tend to come from an engineering or STEM background.

We also shift the focus of the analysis from overall indicators of corporate performance to a specific performance outcome that has received little to no attention so far – namely, cost efficiency. Our findings on the efficiency gains from female representation relate to several recent studies. A large literature has shown that women are socialized into having preferences that are often systematically different from men's (Croson and Gneezy 2009; Dittrich and Leipold 2014). Arano et al. (2010) and Niederle (2017) document gender differences in risk aversion, which may account for corresponding differences in labour market outcomes. Sah et al. (2022) find that the gender gap in risk propensity can also be observed in a sample of CEOs, while Chen et al. (2022) report higher risk-taking behaviour amongst Chinese firms headquartered in counties with a higher male-to-female sex ratio.[3]

A parallel literature has also shown that women have a greater preference for rule-following and ethical behaviour than men. In Italy and China, female bureaucrats are less likely to be investigated and arrested for corruption than their male counterparts, potentially because they tend to 'act more "defensively" in administering their duties' (Decarolis et al. 2023). Female executives have also been shown to reduce the likelihood of firms engaging in financial fraud (Cumming et al. 2015), and to promote the firm's engagement in charitable and socially responsible initiatives (Hafsi and Turgut 2013; Post et al. 2011; Bear et al. 2010; Webb 2004).

Women may bring their gender-specific preferences into the boardroom. Gul et al. (2008) show that companies with gender-diverse boards choose more specialist auditors and demand greater audit effort, as measured by audit fees. They attribute this relationship to a female preference for ethical compliance and to women's higher aversion to risk compared to men's. Using a sample of US firms, Adams and Ferreira (2009) argue that director attendance rates improve with an increase in board gender diversity, and that female directors are five percentage-point more likely to sit on monitoring-related committees than male directors. Similarly, Green and Homroy (2018) find that female directors in Europe are almost ten percentage-points more likely than their male counterparts to sit on the board's audit committees, which is typically responsible for appointing auditors and monitoring the firm's internal financial performance.

Gul et al. (2011) find that female representation on corporate boards improves stock-price informativeness through increased transparency and public financial disclosure. Similarly, Adams and Ferreira (2009) argue that female directors are more likely than their male counterparts to monitor and hold CEOs accountable for poor stock-price performance. Aktaş et al. (2023) show that the 2011 gender quota law in France prompted affected firms to dis-invest in foreign subsidiaries. They, too, attribute this effect to an increase in managerial monitoring, which has the effect of keeping 'corporate empire-building' by (typically, male) CEOs in check (ibid.).

Women may also improve managerial monitoring by contributing distinct types of expertise that are missing in incumbent (male-dominated) boards. Kim and Starks (2016), for instance, find that women directors in the US are more likely to possess functional expertise in the area of risk management, corporate governance and regulation (amongst others). These skills may contribute to more intense monitoring of management by the board, as well as better auditing of the firm's finances.

[3] Whether gender-specific risk preferences extend 'beyond the glass ceiling' is now more controversial. Adams and Ragunathan (2015), for instance, find that banks with more gender-diverse boards do not necessarily have less risk than other banks.

In this paper, we suggest that a potential consequence of tougher managerial monitoring by gender-diverse boards is an increase in the efficiency of the firm's existing operations. In particular, it is plausible that the firms that exert or demand a higher audit effort may as a result become more efficient cost-minimizers. Owing to their (gender-specific) preference for tougher monitoring, women directors may thus improve the ability of the firm to pursue the economic goal of cost minimization. This paper contributes to the literature on gender and corporate governance by examining this possibility.

3 Data

We use data from the restricted-access version of the *Mining Intelligence* dataset published by Glacier Rig Ltd, a Canadian data-analytics company that provides consulting services to the global mining industry. This source is composed of two parts. The 'companies dataset' contains information on the financial performance and board composition of global mining companies. The 'properties dataset' contains information about individual mines – their ownership structure, technical and geographical characteristics, and accounting data such as ore output and production costs.

By matching individual mines with the companies that control them, we constructed an unbalanced panel dataset covering the worldwide gold-mining sector during 2012-2020. We focus specifically on the gold sector for two reasons: i) Individual mine efficiency can be estimated more reliably in a stochastic frontier framework by using a sample of productive units (mines) that operate a similar technology of production; ii) The information available from *Mining Intelligence* is richest and most complete for the gold sector.

The full sample is composed of 136 publicly listed mining companies from 16 headquarter countries. These companies are matched with 294 gold mines across 46 mining countries. While different companies may own stakes in the same mine, we consider the owner of a mine to be the company holding the largest interest share in the mine.[4] The board of the mine's majority shareholder may be assumed to have the greatest influence on the mine's operations and performance. Around 40 percent of the sampled mines for which information on both mine and company location is available are owned by a domestic company headquartered in the same country, while 60 percent are owned by a foreign company. On average, a mining company in our sample owns and operates 2.2 gold mines.

[4]In our full sample, the largest single interest shares in a mine ranges between 25 and 100 percent, with a mean (median) of 94.1 (100) percent.

3.1 Mine-level data

At the mine level, we have information on costs of production, output quantity, a number of (time-invariant) characteristics of the mine (e.g., the ore grade), and the mine's geographical location. This data allows us to specify a rich cost function at the mine level.

[Table 1]

Descriptive statistics on mine-level costs and output are presented in Table 1, Panel A.[5] Following standard industry practice, we make use of the main headline metrics used by gold-mining firms to report costs (World Gold Council, 2021). All-in sustaining costs (AISC) are intended to reflect the full costs of keeping a mine in business. AISC is the sum of two components – total cash costs (C1) and sustaining-capital costs (SC). C1 correspond to short-term production costs (COGS), including those arising from ore extraction and basic processing, and from running the mine site. SC include expenditures intended to keep the mine profitable in the long run (Yapo and Camm 2017; O'Connor et al. 2016). These include exploration costs; the replacement of machinery and equipment; capital expenditures related to safety, health, and the local environment; and costs incurred for mine-site reclamation and rehabilitation.

Actual reporting practice varies significantly from company to company. In the analysis, we focus on the mines owned by companies that report all the cost metrics explained above. For this reason, and because of missing information on the company-level variables, the sample available for estimation is smaller ($N = 352$). In this sample, the total reported costs of production are generally higher than in the full sample, but so is the level of gold output (Table 1, Panel A).

According to the *Mining Technology* magazine, there were 1322 gold mines in operation globally in 2023.[6] The world's largest gold producers are China, Russia, and Australia, followed by Canada and the US. At least for advanced economies such as Australia, Canada and the US, the mines sample obtained from the *Mining Intelligence* dataset is representative of these countries' gold sector.[7] For instance, in 2023 there were 127 active gold mines located in the US, 22 of which (17.3 percent) appear in our full sample. The largest shares of observations in the sample pertain to mines located in Australia (14.3 percent), South Africa (14), the US (7.8) and Canada (6.9). Other important gold producers such as China (1.2) and Russia (4.7) are relatively less well-represented.

3.2 Company-level data

From the *Mining Intelligence* 'companies dataset', we obtained information on the board composition of mining companies, in addition to income-statement, balance-sheet, and firm-demographic information. To measure female

[5] Descriptive statistics on mine and geographic characteristics are not reported in full but available upon request.
[6] See URL: https://www.mining-technology.com/marketdata/five-largest-gold-mines-the-us/
[7] Personal communication with *Mining Intelligence.*

board representation, we divided the number of active female directors by the total number of current directors sitting on the company's board. The share of female directors is intended to capture the influence that women exert on decision-making at the board level − for instance, the intensity of managerial monitoring. In the robustness analysis, we also use indicator variables for companies with at least one (or more than one) female director.

[Figure 1]

Descriptive statistics on board characteristics, specifically those related to gender, are presented in Panel B of Table 1.[8] There is substantial variation in the share of female board members across companies and over time (s.d. = 0.126), with a mean (median) share of 14 percent (12.5 percent). 69 percent of mine-year observations are matched with companies having at least one female director on the board, while around 50 percent are matched with companies with more than one female director. A standard t-test cannot reject the null that the full and estimation samples are drawn from populations with the same mean level of female representation across these three indicators.

In the analysis, we exploit the variation in female boardroom representation *within* companies over time. The within standard deviation in the share of female directors is 0.086, which is only slightly smaller than the between standard deviation (0.103). Figure 1 shows the evolution of female boardroom representation over time, averaging across companies, in the four countries hosting the largest number of gold-producing companies in our sample (Canada, South Africa, Australia, US). The diagrams indicate a clear upward trend in female participation across the four countries. Globally, female participation in mining companies' boards increased from exactly 0 in 2012 to an average of 30.4 percent in 2020 (based on our full sample). In the analysis, we focus specifically on a component of time variation in female board representation that we can plausibly consider to be exogenous to operational efficiency.

At the board-level, we have information on when directors took office, the board size, the age of directors, and whether directors also serve as executives in the company's C-suite or are independent directors. Nearly 6 percent of directors in our sample are new directors, meaning that they took office in the year of observation. On average, the boards of gold-mining companies are composed of 11 directors with 60 years of age. A quarter of directors in the typical board are independent. The share of female directors is quite highly correlated with board size (0.46) and with the share of independent directors (0.61), suggesting that female directors are typically recruited in addition to, rather than as a replacement for, incumbent male directors, and from outside the company rather than from its C-suite. Indeed, both the average board size and its independence increased dramatically in our sample during 2012-2020. In the analysis, we show that our results are robust to controlling for these board

[8]Descriptive statistics on accounting and firm-demographic variables are not reported in full but available upon request.

characteristics.

4 Empirical methods and estimation results

4.1 Estimation strategy

Our analysis proceeds in two steps, as outlined by Greene (2012). First, we use a stochastic cost frontier model to estimate the efficiency level of individual mines. Second, we examine the effects of female board representation, measured at the company-level, on the indicators of mine-level inefficiency obtained in the first step.

First Step. Using our mine-level data, we estimate the following cost function:

$$\ln C_{irt} = \alpha \ln Q_{irt} + (\rho_r + \theta_r t + \phi_r t^2) + \beta X_{ir} + \eta_{irt} \tag{1a}$$

where i indexes' mines located in sub-national regions r (as defined in the *Mining Intelligence* dataset), and t refers to time. C_{irt} represents any one of the three mine-level cost metrics described in section 3 (AISC, C1 and SC). Q_{irt} denotes the quantity of output, measured in MID (metal in doré) units, a standard metric used in the gold industry. The terms $\rho_r + \theta_r t + \phi_r t^2$ model region-specific, non-monotonic time trends that may arise from price variations in local labour and input markets. This term allows wages and other input prices to vary both across regions and over time.

All the gold mines in our sample are industrial (as opposed to artisanal) operations and may be assumed to use the same gold-mining technology. Nevertheless, our specification includes X_{ij}, a rich set of geographical characteristics that may be associated with local technology adaptations. These include: latitude, longitude, a measure of the site's remoteness (in logs)[9], the mine's ore grade (in logs), extant reserves (in logs), and the mine type (e.g., underground vs. open-pit). Given the short time span, we assume no technical change.[10]

Since the theoretical cost function represents an ideal − the frontier of minimum costs that can be feasibly achieved with a given technology in a given price environment − any deviation from it can be interpreted as arising from individual inefficiencies. Thus, the error term in equation (1) is allowed to have a composite form, which follows Battese and Coelli's (1992) 'normal-truncated normal' specification:[11]

$$\eta_{irt} = e_{irt} + u_{irt} \tag{1b}$$

e_{irt} is the standard idiosyncratic disturbance reflecting model error and random-sampling variability, and is assumed to have a *symmetric* (normal) distribution. u_{irt} denotes the so-called 'inefficiency term' − a *one-sided*,

[9] Defined as the travel distance to the nearest urban center.
[10] To the extent that technical change takes place and is trended, it will be captured by the region-specific trend terms.
[11] This specification is also known as the 'time-varying decay' model.

time-varying residual capturing upward deviations from the minimum-cost frontier. When $u_{irt} = 0$, the mine is operating its technology at full efficiency at time t. When $u_{irt} > 0$, the mine is spending more than it could given its technology and the market environment it faces. u_{irt} is assumed to have the following truncated-normal form, where N^T refers to the truncated-normal distribution and η is a constant:

$$u_{irt} = u_{ir} \times \exp\left[-\eta(t - T_{ir})\right] \tag{2a}$$

$$u_{ir} \sim N^T(0, \sigma_u^2) \tag{2b}$$

We first estimate the parameters of the cost function (eq. 1a) by maximum likelihood (ML). Next, we use the method of Jondrow et al. (1982) to obtain estimates of u_{irt} for each of the three cost variables used on the left-hand side of equation (1a) $-$ that is AISC, C1 and SC.[12] The resulting inefficiency terms have the same unit of measurement as the corresponding dependent (cost) variables (log of mln US$).

[Table 2]

Table 2 shows the estimates of equation (1a) for the three cost metrics. As expected, higher levels of production increase total costs with an elasticity ranging between 0.8 and 1, indicating increasing returns to scale.[13] Except for model (2), both the region-level trends and the set of mine-level controls (X_{ir}) enter the cost function as jointly significant, suggesting time variation in local input prices and/or technology adaptation across locations.[14]

[Table 3]

Table 3 shows descriptive statistics for the inefficiency terms (u_{irt}) obtained in post-estimation, which we call AISC-, C1- and SC-inefficiencies, respectively. The table distinguishes between the full sample and the sample available for estimation in the second step of our analysis ($N = 352$). Across the three measures of cost-inefficiency, t-tests cannot reject the null (at the 5 percent level, at least) that the full and estimation samples

[12] To extract the inefficiency term u_{irt} from the composite error term (η_{irt}), Jondrow et al. (1982) propose to exploit the assumed conditional distribution of u_{irt} given $\hat{\eta}_{irt}$. The estimates of individual inefficiencies can then be obtained using the mean $E(u_{irt}|\hat{\eta}_{irt})$ of the conditional distribution. In alternative specifications reported in the Appendix, we also test the robustness of our results to using an alternative method proposed by Battese and Coelli (1988), which uses $E[\exp(-u_{irt}|\hat{\eta}_{irt})]$ to recover the estimates of u_{irt}.

[13] Increasing returns to scale are in line with previous findings from the mining sector (e.g., Boyd 1987).

[14] The OLS estimates of the cost-function parameters (eq. 1a), which we obtained for comparison, are fairly similar to the ML estimates of the cost-frontier model allowing for a composite error term. For example, the OLS estimates of α for models (1)-(3) are, respectively, 0.841 (s.e. = 0.047), 0.819 (0.035), 1.014 (0.079).

are drawn from populations with the same mean level of inefficiency. Together with previously reported t-tests showing no difference in mean female participation across the full and estimation samples (Table 1), these findings mitigate the concern that our main results may be an artefact of sample selection.

Both C1- (0.51) and SC-inefficiency (0.40) are moderately highly correlated with AISC-inefficiency, but less so with each other (0.13), confirming that these two cost metrics convey complementary information. For illustration, Figure 2 plots the distribution of AISC-inefficiency for the sampled mines located in Russia, Canada and the United States, respectively. The right-skewed, truncated-normal distribution of this variable is apparent in the diagram. It is also evident that the mines located in countries with a comparative advantage in resource extraction (Russia and, to a lesser extent, Canada) are, on average, less inefficient than the mines located in a more diversified economy such as the United States.

[Figure 2]

Figure 3 shows the evolution of AISC-inefficiency over time, averaging across mines, in the four countries hosting the largest number of gold mines in our estimation sample, namely South Africa, Canada, the United States and Australia. Across these countries, AISC-inefficiency displays an upward trend, which could be driven by changes in industry-specific regulation. This finding underscores the need to control for global trends in the second-step of the empirical analysis.

[Figure 3]

Second step. To examine the impact of female boardroom representation on mine-level inefficiency, we begin by estimating the following benchmark OLS model:

$$u_{icjt} = aFem_{jt} + \omega_c + \sigma_j + \tau_t + bX_{jt} + \epsilon_{icjt} \tag{3}$$

Here, we match the mine-level with the company-level information. Thus, u_{icjt} denotes the AISC-, C1- or SC-inefficiency of mine i at time t, where the mine is located in country c and is controlled by company j at time t. Fem_{jt} is the share of female directors sitting on the board of company j at time t. ω_c is a full set of mining-country fixed effects, capturing the country-level differences in mining-sector efficiency illustrated in Figure 2. τ_t are time effects, which control flexibly for global trends in female representation and inefficiency levels. σ_j denotes company fixed-effects, which control for company-specific unobservable characteristics (e.g., management culture, corporate ethics) that may influence both female representation and cost efficiency. Lastly, X_{jt} is a vector of additional time-varying company characteristics that may affect cost-efficiency, potentially confounding the OLS estimate of

a, our parameter of interest. Here, we consider a measure of firm size (the log of total assets) and indebtedness (debt to equity ratio).[15]

ϵ_{icjt}, the error term, reflects idiosyncratic variation in cost-inefficiency across mines, companies and time. Even after controlling for σ_j, some of this variation is likely to be common to mines belonging to the same company. Thus, following MacKinnon et al.'s (2023, 278) recommendation, we cluster the OLS standard errors conservatively at the coarsest feasible level (the company), allowing for general patterns of residual correlation both across mines owned by the same company and within the same company over time.[16] In section 4.5, we test the robustness of our conclusions to alternative clustering structures. To increase the efficiency of the OLS estimator, we also weight the observations by the interest share (%) held by company j in mine i at time t. This procedure has the effect of reducing the influence of observations referring to mines that are not wholly owned by the company to which they are matched in the dataset.[17]

4.2 OLS results

The estimates of equation (3) for our three different cost-inefficiency outcomes are presented in Table 4. The models in column 1, which are presented for comparison, control for mine-country (ω_c) and time-period (τ_t) effects only. The estimates of coefficient a are negative across Panels A-C, indicating that a higher share of female directors on a company's board is associated with lower inefficiency across the mines controlled by the company. Yet, the estimates are always statistically insignificant.

[Table 4]

The models in column 2, which we consider as our baseline specifications, add company fixed effects (σ_j) to the regression equation. The models in column 3 also add the full set of company-level controls (X_{jt}). These specifications only use variation over time in female board representation within companies for identification. The estimates of coefficient a are now negative and statistically significant at conventional levels across Panels A-C. For the average mining company, a one standard deviation increase (12 percentage points) in the share of female directors is associated with a 0.7, 0.5 and 2 percent reduction in AISC-, C1- and SC-inefficiency, respectively (based on model 2), across the mines controlled by the company. These effects correspond approximately to 4.4,

[15] The potentially confounding influence of other *board* characteristics (as opposed to company-level characteristics such as assets and indebtedness) is considered later in the analysis.

[16] A coarser clustering level could be the country where the company is headquartered. With only seven such countries (in the estimation sample), however, this clustering structure is infeasible, as the small number of clusters would cause the hypothesis tests to over-reject the null.

[17] Our results are qualitatively robust to using an estimation procedure (OLS or 2SLS) without weights, however. The results are available upon request.

2.8 and 8.9 percent of a standard deviation of the three inefficiency indices, respectively. Figure 4 visualizes the partial relationship between the share of female directors and AISC-inefficiency estimated in columns (2) and (3) of Table 4 (Panel A). The overall beneficial effect of female directors appears to come primarily from efficiency improvements in sustainability-related (SC) cost categories (Panel C).[18] Yet, female directors also appear to make a significant contribution to reducing C1-inefficiency (Panel A). Improvements in both C1- and SC-inefficiency explain the resulting decrease in AISC-inefficiency reported in Panel A (Table 4) and depicted in Figure 4.

While the economic significance of these efficiency gains is quite modest, the estimated effects are still larger than those reported in previous studies that find a beneficial effects of female participation on overall corporate performance. For instance, Green and Homroy (2018) show that a one standard deviation increase in female representation on European corporate boards increases ROA by only 2.6 percent of a standard deviation.

[Figure 4]

As we noted earlier, a possible concern with giving these associations a causal interpretation is that mining companies may select female directors based on their level of operating efficiency. It is also possible that prospective female directors may self-select into systematically more efficient companies − for instance, if observable efficiency is taken as a signal for a female-friendly environment. To mitigate these concerns, model 4 conditions the estimates on two lagged measures of corporate performance that are easily observable by prospective female directors – namely, return on assets (ROA) and asset turnover (a standard measure of operating efficiency).[19] If the selection mechanism depends on these two variables, their inclusion should 'kill' the estimated effect of female directors on mine-level inefficiency. Yet, our findings remain qualitatively unchanged (although the coefficients are now less precisely estimated).

In additional tests not reported in full, we also show that although the indices of cost-inefficiency have a right-skewed distribution (see Figure 2), our results are not driven by influential outliers.[20] We also confirm that our linear specification provides a good fit to the data,[21] and that the efficiency-enhancing effects of female directors hold homogenously across companies reporting different levels of profitability.[22]

[18] Still, the effects on SC-inefficiency are generally less precisely estimated

[19] Asset turnover is measured as sales revenues over assets.

[20] To do so, we inspect partial-correlation plots (as in Figure 4) and leverage-versus-squared-residual plots for all the regressions reported in Table 4. In one case, we dropped some influential outliers, which left our results unaltered (full results available upon request).

[21] To check this, we cut the distribution of Fem_{jt} into three terciles and show that the estimated efficiency effects hold uniformly across the distribution (full results available upon request).

[22] To do so, we interact Fem_{jt} with the company's average ROA during the sample period (which may be taken as a measure of the quality of management). The coefficients of the interaction terms are always small and statistically insignificant (full results available upon request).

4.3 IV estimation

The OLS estimates do not identify the causal effect of female directors on cost efficiency if board members are chosen endogenously with the aim of optimizing operating efficiency (Ahern and Dittmar 2012), or if efficient female directors are systematically discriminated against (Comi et al. 2020). We address this concern (and at the same time, any remaining concern related to selection and self-selection) by using instrumental-variable (IV) regressions.[23] In particular, we use an IV to isolate a component of variation in female representation on the board that is plausibly unrelated to the firm's endogenous choice of hiring female directors.

The instrument: We argue that the choice of the board's gender composition is subject to exogeneous peer effects (Bulow et al. 1985; Bustamante and Fresard 2021). Firms may be under pressure to hire more women if their direct competitors embark on efforts to promote gender equality in the boardroom. We suggest that these effects are likely to reflect the horizontal transmission of cultural norms regarding female leadership across peer firms. As such, they are likely to be strongest amongst peers headquartered in the same city. Thus, we define the mean share of female directors amongst the other $[N_t^{(C)} - 1]$ companies headquartered in the same city C as company j at time t, and use this jack-knifed average as an instrument for Fem_{jt}:

$$\overline{Fem}_{jt}^{(C)} = \left(\sum_{k \in C} Fem_{kt} - Fem_{jt} \right) \Big/ \left[N_t^{(C)} - 1 \right] \tag{4}$$

We interpret $\overline{Fem}_{jt}^{(C)}$ as reflecting city-wide patterns of female corporate participation $-$ for example, those driven by cultural norms and cultural change. Note that, similar to the company's own share of female directors (Fem_{jt}), $\overline{Fem}_{jt}^{(C)}$ varies at the firm-year level. In the sample available for estimation, $\overline{Fem}_{jt}^{(C)}$ is highly correlated with Fem_{jt} (0.81).

Our identification assumption is that the average level of female participation amongst a company's peers only affects the cost-efficiency of its operations by influencing the company's choice to hire female directors. Here, we use our data to assess the plausibility of this assumption. Although highly correlated with Fem_{jt}, our instrument is uncorrelated with all the company-level characteristics that we observe $-$ including size, indebtedness, ROA and asset turnover (see Appendix A). It is also uncorrelated with other board-level characteristics such as average directors' age and the share of new directors. The only exceptions are the size and independence of the board, which are highly correlated with the presence of female directors. Thus, the data appear to rule out other potential channels through which city-level female representation could affect a company's cost-efficiency. Nevertheless, in subsequent specifications, we explicitly control for these (and other) board characteristics in the structural equation.

Another possible concern is that an increase in $\overline{Fem}_{jt}^{(C)}$ may induce efficiency effects amongst the company's peers, and thereby influence the company's own efficiency performance through spill-over effects. If so, the 2SLS

[23] None of the countries in which our sampled companies are headquartered introduced gender quotas during 2012-2020. For this reason, we cannot use quotas as instrumental variables.

estimate of parameter a in equation (3) may simply reflect efficiency spill-overs across companies headquartered in the same city, rather than the impact of the company's own female directors. To mitigate this concern, we always include the jack-knifed average ROA of the company's city-level peers as an additional control in the second-stage of the IV estimation.[24]

[Table 5]

Empirical results. 2SLS estimates of equation (3) are reported in Table 5. All models include mine-country, year and company fixed effects, as well as peers' ROA. Column 1 shows the estimated coefficients of an OLS benchmark model. Columns 2 and 3 report 2SLS models based on the city-level instrument, with and without company-level controls (X_{jt}). Panel D reports the first-stage results, which are common to all the three regressions shown across Panels A-C.[25]

In model 2, the 2SLS estimates of parameter a in equation (3) are always negative, statistically significant and almost *twice* as large in absolute magnitude as the corresponding OLS estimates reported in column 1. This finding suggests that the OLS estimates may be subject to downward bias (in absolute terms). A possible explanation is that, in equilibrium, efficiency-enhancing women candidates are subject to some statistical or taste-based discrimination in the executive labour market (see Fernandez-Mateo and Fernandez [2016] for a discussion). Accordingly, an exogenous increase in female participation (as captured by the first-stage fitted values from the 2SLS procedure) is associated with larger efficiency gains than implied by OLS. A one standard deviation increase (12 percentage points) in the share of female directors translates into a 1.9, 1.3 and 3.4 percent decrease in AISC-, C1- and SC-inefficiency, respectively (based on model 2). These effects correspond approximately to 12, 8 and 15 percent of a standard deviation of the three inefficiency indices, respectively − a quantitatively large effect.

The remaining columns of Table 5 test the robustness of these findings to using alternative specifications. The model in column 3 adds two time-varying controls (ROA and indebtedness) to the regression equation, leading to similar results. In columns 4 and 5, we construct alternative instruments for Fem_{jt} by using province- and country-level (jack-knifed) averages of the share of female directors, instead of city-level averages. Since these two instruments are less strongly correlated with the endogenous variable (as evidenced by lower first-stage F-statistics and partial R-squared), the corresponding coefficients on Fem_{ij} are less precisely estimated than those obtained with the city-level instrument. Yet, they remain statistically significant and fairly similar in magnitude to those reported in column 2.

[24] In Appendix A, however, we show that conditional on company fixed effects, $\overline{Fem}_{jt}^{(C)}$ is uncorrelated with peers' ROA. Our results are robust to dropping this control (full results available upon request).

[25] After partialling out the other covariates, $\overline{Fem}_{Ct}$ explains 53-54 percent of within-firm variation in Fem_{jt}. The F-statistic for the test of weak identification is always greater than the relevant Stock-Yogo critical value (16.38), leading to a rejection of the null of under-identification.

[Table 6]

Mechanisms. Why do female directors enhance cost efficiency? We noted previous findings showing that women executives have gender-specific preferences (Sah et al. 2022), and that they tend to allocate more effort to monitoring and auditing (for instance, by sitting on monitoring-related committees) than their male peers (Gul et al. 2008; Adams and Ferreira 2009). While our data do not allow us to perform a direct test of these mechanisms, we are able to rule out several alternative explanations.

Incoming women directors are typically recruited from outside the company, and added onto the existing board. Thus, it is possible that they may promote efficiency-enhancing board decisions by increasing the independence and/or the size of the board, rather than by imposing a gender-specific preference for tougher monitoring. Previous results also showed that women directors tend to be younger than their male counterparts (Ferrari et al. 2020). Thus, they could promote efficiency by bringing in new (but not necessarily gender-specific) values and perspectives.

To assess these possibilities, in Table 6 we run 2SLS regressions (as in Table 5, column 2) that include the variables presented in Table 1, Panel B (board size, the share of independent directors, average directors' age, and the share of new directors) as additional controls in the structural equation. In columns 1-4, these variables are entered individually; in column 5 they are entered simultaneously in the same regression. Our main findings remain qualitatively unchanged throughout. We conclude that the efficiency gains associated with more gender-diverse boards cannot be attributed to other board-level changes (e.g., in size or independence) induced by the entry of women.[26] Rather, the evidence is consistent with these efficiency gains resulting from changes in corporate strategies (e.g., tougher monitoring) induced by the gender-specific preferences of incoming female directors.

4.4 Effects on firm profitability

To complete the analysis, we examine whether the efficiency gains induced by female directors translate into better overall performance outcomes in our sample of companies. To do, we estimate the effects of the share of female directors on various accounting-based measures of firm profitability, including gross income (= gross profits), operating income, pretax income and net income, all defined as a share of total assets. We estimate the following equation at the company level:

[26] All control variables enter as insignificant except for board size, which is found to increase AISC- and C1-inefficiency (full results available upon request). This result is consistent with previous findings of a negative effect of board size on corporate performance (Adams and Ferreira 2009, 305-307).

$$Y_{jt} = aFem_{jt} + \sigma_j + \tau_t + bX_{jt} + \epsilon_{jt} \tag{5}$$

where Y_{jt} refers to any of the above-mentioned performance outcomes and the other symbols are defined as in equation 3.

[Table 7]

The 2SLS estimates of the parameters in equation (5), using the city-level instrument for Fem_{jt}, are presented in Table 7.[27] The effects of gender-diverse boards on the ratio of gross and operating income to assets are small and insignificant, while the effects on the ratio of pre-tax and net income to assets are positive but very imprecisely estimated. Although these company-level findings should be interpreted with caution, we conclude that female directors may increase the efficiency of a company's operations without necessarily improving the overall performance of the company as measured by (short-term) accounting profitability. This finding resonates with those of Adams and Ferreira (2009), who link female directors to increased monitoring, but potentially lower profits and market value, and is generally in line with the balance of evidence in the literature.

4.5 Robustness analysis

Here, we conduct several additional robustness tests on the IV results reported in Table 5. In Appendix B, we test the sensitivity of our findings to the specification of the stochastic frontier model. We show that the 2SLS parameter estimates are qualitatively robust to: using Battese and Coelli's (1988) method of post-estimating the inefficiency term, instead of Jondrow et al.'s (1982); removing the technology controls (X_i) from equation (1); using linear, instead of quadratic, trends ($\rho_r + \theta_r t$); and adding a quadratic term in output to the cost function (as in the trans-log functional form). These findings indicate that our conclusions are not an artifact of the specific stochastic-frontier specification used in step one of the main analysis. They also demonstrate the robustness of our findings to using alternative measures of cost-inefficiency as dependent variables in equation (3).

In Appendix C, by contrast, we investigate the robustness of our results to using alternative independent variables. Specifically, we use the two binary indicators of female boardroom representation (at least one female director on the board, more than one female director) summarized in Table 1, Panel B, instead of the share of female directors. Our findings remain qualitatively unaltered. The companies with at least (more than) one female director have mines that are around 12 (13) percent more cost-efficient, using the AISC metric, than companies without (or with only one) female director. Our conclusions do not depend critically on one particular definition of female representation.

[27] All regressions control for the jack-knifed average of ROA amongst the own company's city peers.

In Appendix D, we demonstrate the robustness of our findings to alternative specifications of equation (3). We show that our results remain qualitatively unchanged when the mine-country FE (ω_c) are omitted, or replaced with mine-continent FE. The results are also robust to replacing company FE (σ_j) with company-city FE. Lastly, the results do not change substantially when the potential influence of trended unobservables is controlled for by including a full set of company-specific linear trend terms ($h_j + k_j t$) in the regressions .

Lastly, we test the robustness of our conclusions to the choice of inference procedure and clustering structure. Throughout the analysis we assumed that in large samples the 2SLS z-statistics follow an approximately normal distribution. It is now well known that when the number of clusters is small (< 42) asymptotic theory may provide a poor guide to the distribution of test statistics even in large samples. An increasingly popular approach is to base inference on an empirical bootstrap distribution obtained by resampling the data cluster by cluster.[28] In Appendix E, we reproduce the z-statistics and asymptotic p-values for the estimates shown in Table 5, column 2. We then compare these p-values to those obtained using the wild restricted efficient (WRE) bootstrap procedure for instrumental-variable models (Davidson and MacKinnon (2010). We report variants that bootstrap either the z-statistic or only its numerator (the 'bootstrap-c' procedure advocated by Young [2022]). For AISC- and C1-inefficiency (but not for SC-inefficiency), hypothesis tests based on bootstrap critical values always reject the null of no effect, confirming the validity of inference in our main findings.

Another concern is that the model disturbances may be correlated across mines located in the same country, in addition to being correlated across mines owned by the same company. Neither dimension of intra-cluster correlation (company and mine country) is eliminated entirely by including a corresponding set of cluster fixed effects, as in equation 3. Thus, in Appendix E, we also two-way cluster the 2SLS standard errors by company and mine country (as in Cameron et al., [2011]). We then perform hypothesis tests using either asymptotic or bootstrap critical values. In addition, we also two-way cluster the standard errors by company and time, allowing for residual inefficiencies to be correlated across mines owned by different companies in the same year. In all cases, our main conclusions remain qualitatively unchanged.

5 Conclusion

Previous studies found that women directors exert a higher monitoring and audit effort than their male peers (Gul et al. 2008; Adams and Ferreira 2009; Green and Homroy 2018; Nekhili et al. 2020), potentially because of gender-specific preferences for low risk and ethical compliance. We suggested that a possible consequence of tougher monitoring by gender-diverse boards may be an improvement in the firm's operational efficiency.

Using detailed mine-level data from a global representative sample of gold mines, we present the first evidence

[28] The p-values are then calculated as the proportion of bootstrap z-statistics that are larger than the z-statistic obtained from the original sample. MacKinnon et al. (2023, 297) recommend reporting several variants of bootstrap p-values 'as a matter of course'.

on the impact of female directors on cost-efficiency. Using a stochastic frontier methodology and instrumental-variable regressions, we find that an increase in female representation on the parent company's board translates into sizeable cost-efficiency gains in the mining operations controlled by the parent company. These effects are not mediated by other changes in board characteristics induced by the entry of women (e.g., an increase in board independence). Rather, they appear to relate specifically to the gender of incoming directors. While our data do not allow us to test the mechanism directly, a plausible explanation that is in line with previous findings in the literature is that women directors have a gender-specific preference for tougher monitoring of the company's internal finances.

Yet, we find no evidence in our sample that the efficiency gains generated by women directors translate into higher profitability. Our evidence is supportive of the view that women directors bring a distinctly female style of corporate leadership to the board, without necessarily having a systematic effect on the company's overall performance, in line with previous findings in the literature (Johnson et al. 2013; Post and Byron 2015).

The findings of this study have important policy and managerial implications, particularly in the context of gender diversity and leadership styles. The observed impact of female directors on cost-efficiency in mining operations suggests that policies promoting gender diversity in corporate boards can contribute to improving a firm's operations. In particular, decision-makers may consider recruiting female directors when the organization's priority is operational efficiency. Yet, they should recognize that the efficiency gains catalysed by greater female participation, while substantial, may not necessarily translate into higher profitability. This insight emphasizes the need for a nuanced understanding of the role of women in leadership. Decision-makers should consider crafting initiatives that not only encourage gender diversity but also foster an appreciation (and tolerance) for diverse leadership styles. This approach challenges existing norms and expectations about women in top management, paving the way for a more inclusive and open-minded corporate culture. By acknowledging and valuing the distinctive contributions of female leaders, policies and strategies can be designed to create environments that truly harness the benefits of diverse perspectives within corporate decision-making bodies.

While our data only allow us to make inference about the population of gold-mining firms, our findings are arguably relevant for the mining industry more generally, and potentially for other industries, too. Future research should examine alternative dimensions of efficiency (e.g., technical or profit efficiency), and further elucidate the causal channels linking female participation to efficiency outcomes.

6 References

Abraham L (2023) The gender gap in performance reviews. *J. Econ. Behav. Org.* 214: 459-492.

Adams .B, Ferreira D (2009) Women in the Boardroom and Their Impact on Governance and Performance. *J. Fin. Econ.* 94: 291-309

Adams RB, Ragunathan V (2015) Lehman Sisters. FIRM Research paper, 3 September 2015

Ahern KR, Dittmar AK (2012) The Changing of the Boards: The Impact on Firm Valuation of Mandated Female Board Representation. *Quart. J. Econ.* 127: 137-197.

Aktaş K, Gattai V, Natale P (2023) Board gender quotas and outward foreign direct investment: Evidence from France. *Can. J. Econ.* 56(4): 1291-1321

Arano K, Parker C, Terry RL (2010) Gender-based risk aversion and retirement asset allocation. *Econ. Inq.* 48: 147-155

Battese GE, Coelli TJ, (1988) Prediction of firm-level technical efficiencies with a generalized frontier production function and panel data. *J. Economet.* 38: 387-399

Battese GE, Coelli TJ (1992) Frontier production functions, technical efficiency and panel data: With application to paddy farmers in India. *J. Prod. An.* 3: 153-169

Bear S, Rahman N, Post C (2010) The impact of board diversity and gender composition on corporate social responsibility and firm reputation. (J. Bus. Eth.) 97: 207–221

Bennouri M, Chtioui T, Nagati H, Nekhili M (2018) Female board directorship and firm performance: What really matters? *J. Bank. Fin.* 88: 267-291

Boyd GA (1987) Factor Intensity and Site Geology as Determinants of Returns to Scale in Coal Mining. *Rev. Econ. Stat.* 69: 18-23

Bulow JI, Geanakopulos JD, Klemperer PD (1985) Multimarket oligopoly: Strategic substitutes and complements. *J. Pol. Econ.* 93(3): 488–511.

Bustamante MC, Frésard L (2021) Does firm investment respond to peers' investment? *Manag. Sci.* 67(8): 4703–4

Cameron AC, Gelbach JB, Miller DL (2011) Robust inference with multiway clustering. *J. Bus. Econom. Statist.* 29: 238-249

Carbonero F, Devicienti F, Manello A, Vannoni D (2021) Women on board and firm export attitudes: Evidence from Italy *J. Econ. Behav. Org.* 192: 159-175

Chen Z, Huang X, Zhang L (2022) Local gender imbalance and corporate risk-taking. *J. Econ. Behav. Org.* 198: 650-672

Comi S, Grasseni M, Origo F, Pagani L (2020) Where women make a difference: Gender quotas and firm's performance in three European countries. *ILR Rev.* 73: 1-26

Croson R, Gneezy U (2009) Gender differences in preferences. *J. Econ. Litt.* 47 (2): 448-74.

Cumming D, Leung TY, Rui O (2015) Gender diversity and securities fraud. *Acad. Manag. J.* 58: 1572–1593

Davidson R, MacKinnon JG (2010) Wild bootstrap tests for IV regression. *J. Bus. Econom. Stat.* 28(1): 128-44.

Decarolis F, Fisman R, Pinotti P, Vannutelli S, Wang Y (2023) Gender and bureaucratic corruption: Evidence from two countries. *J. Law Econ. Org.* 39: 557-585

Dittrich M, Leipold K (2014) Gender differences in time preferences. *Econ. Lett.* 122: 413-415

Fernandez-Mateo I, Fernandez RM (2016) Bending the pipeline? Executive search and gender inequality in hiring for top management. *Manag. Sci.* 62(12): 3636-3655

Ferrari G, Ferraro V, Profeta P, Pronzato C (2022) Do board gender quotas matter? Selection, performance and stock market efffects. *Manag. Sci.* 68: 5618-5643

Green CP, Homroy S (2018) Female directors, board committees and firm performance. *Eur. Econ. Rev.* 102: 19-38

Greene WH (2012) *Econometric Analysis.* 7th ed. Upper Saddle River, NJ: Prentice Hall

Gregory-Smith I, Main BGM, O'Reilly CA III (2014) Appointments, Pay and Performance in UK Boardrooms by Gender. *Econ. J.* 124: F109-F128

Gul F, Srinidhi B, Tsui JSL (2008) Board diversity and the demand for higher audit effort. Unpublished paper, SSRN 1359450

Gul F, Srinidhi B, Ng AC (2011) Does board gender diversity improve the informativeness of stock prices? *J. Acc. Econ.* 51: 314-338

Hafsi T, Turgut G (2013) Boardroom Diversity and its Effect on Social Performance: Conceptualization and Empirical Evidence. *J. Bus. Eth.* 112: 463–479

Harrison DA, Klein KJ (2007) What's the Difference? Diversity Constructs as Separation, Variety, or Disparity in Organizations. *Acad. Manag. Rev.* 32(4): 1199-1228

Johnson SG, Schnatterly K, Hill AD (2013) Board composition beyond independence: Social capital, human capital and demographics. *J. Manag.* 39(1): 232-262

Jondrow J, Lovell CAK, Materov IS, Schmidt P (1982) On the estimation of technical inefficiency in the stochastic frontier production function model. *J. Economet.* 19: 233–238.

Kim D, Starks LT (2016) Gender diversity on corporate boards: Do women contribute unique skills? *Am. Econ. Rev. P.&P.* 106: 267-271

Matsa DA, Miller AR (2013) A Female Style in Corporate Leadership? Evidence from Quotas. *Am. Econ. J.: Applied* 5(3): 136-69.

MacKinnon JG, Nielsen MO, Webb MD (2023) Cluster-robust inference: A guide to empirical practice. *J. Econom.* 232: 272-299

Nekhili M, Gull AA, Chtioui T, Radhouane I(2020) Gender-diverse boards and audit fees: What difference does gender quota legislation make? *J. Bus. Fin. Acc.* 47: 52-99

Niederle M (2017) A gender agenda: a progress report on competitiveness. *Am. Econ. Rev.* 107: 115-119

O'Connor FA, Lucey BM, Baur DG (2016) Do gold prices cause production costs? International evidence from country and company data. *J. Int. Fin. Mark. Inst. Money.* 40: 186-196

Post C, Byron K (2015) Women on boards and firm financial performance: A meta-analysis. *Acad. Manag. Rev.* 58(5): 1546–1571

Post C, Rahman N, Rubow E (2011) Green Governance: Boards of Directors' Composition and Environmental Corporate Social Responsibility. *Bus. Soc.* 50: 189-213

Sah NB, Adhikari HP, Krolikowski MW, Malm J, Nguyen TT (2022) CEO gender and risk aversion: Further evidence using the composition of firm's cash. *J. Behav. Exp. Fin.*, 100595

Webb E (2004) An Examination of Socially Responsible Firms' Board Structure. *J. Manag. Gov.* 8: 255-277.

Keller W, Molina T, Olney WW (2023) The Gender Gap Among Top Business Executives. *J. Econ. Behav. & Org.* 211: 270-286.

Yapo AG, Camm TW (2017) All-in sustaining cost analysis: Pros and cons. Paper presented at the SME Annual Meeting, Denver, CO, February 19-22

World Gold Council 2021. *Methodology: Production costs data set*, available at: https://www.gold.org/download/file/13470/Production$_{c}$*osts*$_{m}$*ethodology.pdf*

TABLE 1 - Descriptive statistics

	Full sample				Estimation sample			
	N	mean	(s.d.)	{min, max}	N	mean	(s.d.)	[p-value]
Panel A: Costs of production and gold output (mine level)								
AISC: All-in Sustaining Costs (mln US$)	986	215	(201)	{0.9, 1739}	352	230	(213)	[0.253]
C1: Cash costs (mln US$)	1351	125.5	(139.5)	{0.3, 1627}	352	166	(154)	[0.000]
SC: Sustaining-capital costs (mln US$)	605	54.4	(74.9)	{−344, 484}	352	63.9	(70.3)	[0.053]
Gold output: Metal in doré (kozt)	2092	209	(232)	{0.2, 2110}	352	256	(255)	[0.001]
Panel B: Boards of directors (company level)								
Share female directors (fraction)	818	0.143	(0.126)	{0, 0.5}	352	0.139	(0.121)	[0.614]
At least one female director (dummy)	818	0.688	(0.464)	{0, 1}	352	0.690	(0.463)	[0.944]
More than one female director (dummy)	818	0.494	(0.500)	{0, 1}	352	0.514	(0.501)	[0.524]
Board size (n)	818	10.8	(7.2)	{1, 47}	352	11.8	(7.8)	[0.034]
Share independent directors (fraction)	818	0.249	(0.215)	{0, 1}	352	0.257	(0.224)	[0.565]
Average board age (years)	686	60.3	(4.7)	{42, 77}	296	60.6	(4.6)	[0.538]
Share new directors (fraction)	890	0.058	(0.104)	{0, 0.75}	352	0.075	(0.128)	[0.015]

Notes: AISC is equal to the sum of C1 and SC. The p-value is for a test of the equality of the means across the full and estimation samples (assuming independent samples). N stands for number of observations. 'kozt' stands for thousand troy ounces (a standard quantity metric in the gold industry).

FIGURE 1 - Average share of female board directors in selected company-HQ countries

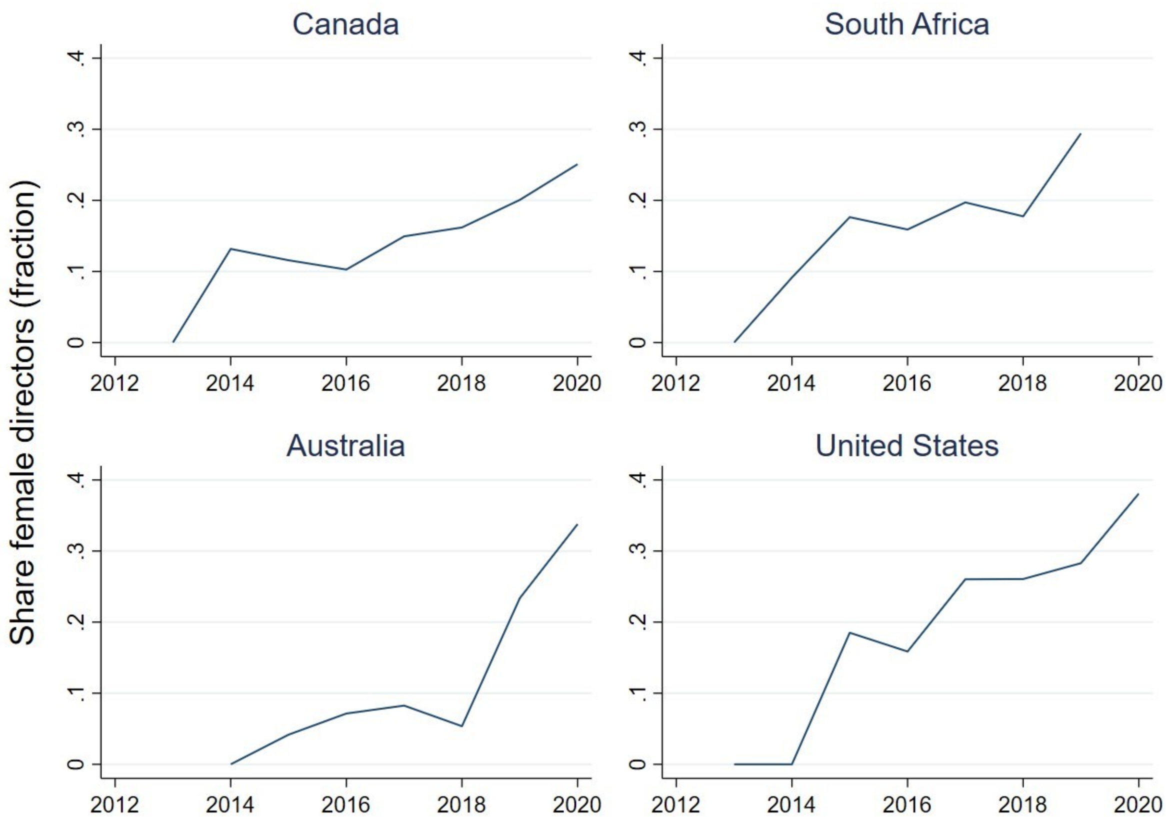

Notes: The four countries are where the largest number of gold-producing companies are headquartered in our estimation sample. Collectively, these four countries account for 88 percent of observations in the estimation sample.

TABLE 2 - Stochastic Cost Frontier Analysis

Dependent variable:	ln(AISC) (1)	ln(C1) (2)	ln(SC) (3)
ln(Gold output)	0.824***	0.800***	0.990***
	(0.039)	(0.041)	(0.069)
Province-level quadratic trends	YES	YES	YES
Joint test [p-value]	[0.000]	[0.000]	[0.000]
Mine-level controls	YES	YES	YES
Joint test [p-value]	[0.000]	[0.527]	[0.001]
σ_u	5.121	4.281	5.990
σ_e	0.150	0.034	0.336
Mines	184	201	130
Observations	916	1251	557

Notes: ML regressions with robust standard errors clustered at the mine level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Gold output is measured as metal in doré (MID). The province-level trends control for prices in local input markets. The mine-level controls include: latitude, longitude, the log distance from the nearest urban settlement (remoteness), the mine deposits grade (in logs), the mine's gold reserves (in logs), and dummies for different mine types (open-pit, underground, mixed, tailings mine).

TABLE 3 - Mine-level cost inefficiencies: Descriptive statistics

	N	mean	(s.d.)	{min, max}	[p-value]
AISC-inefficiency					
Full sample	916	0.145	(0.181)	{0.021, 1.619}	
Estimation sample	352	0.131	(0.165)	{0.022, 0.905}	[0.207]
C1-inefficiency					
Full sample	1251	0.132	(0.149)	{0.015, 0.974}	
Estimation sample	352	0.150	(0.185)	{0.015, 0.974}	[0.059]
SC-inefficiency					
Full sample	557	0.178	(0.211)	{0.030, 2.227}	
Estimation sample	352	0.178	(0.225)	{0.031, 2.227}	[0.983]

Notes: The p-value is for a test of the equality of the means across the full and estimation samples (assuming independent samples). N stands for number of observations.

FIGURE 2 - Distribution of AISC-inefficiency in three mining countries

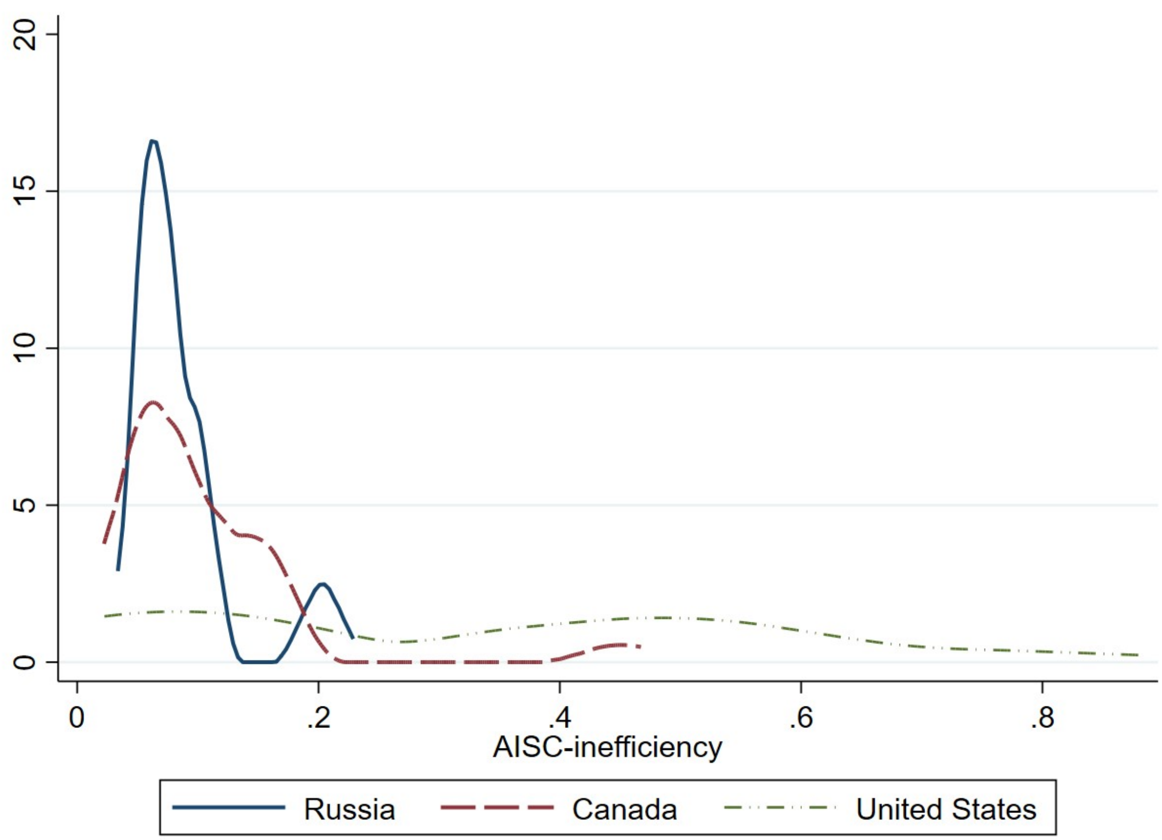

Notes: The kernel density plots display the distribution of AISC-inefficiency for the sampled mines located in Russia, Canada and the United States.

FIGURE 3 - The evolution of AISC-inefficiency over time in selected mining countries

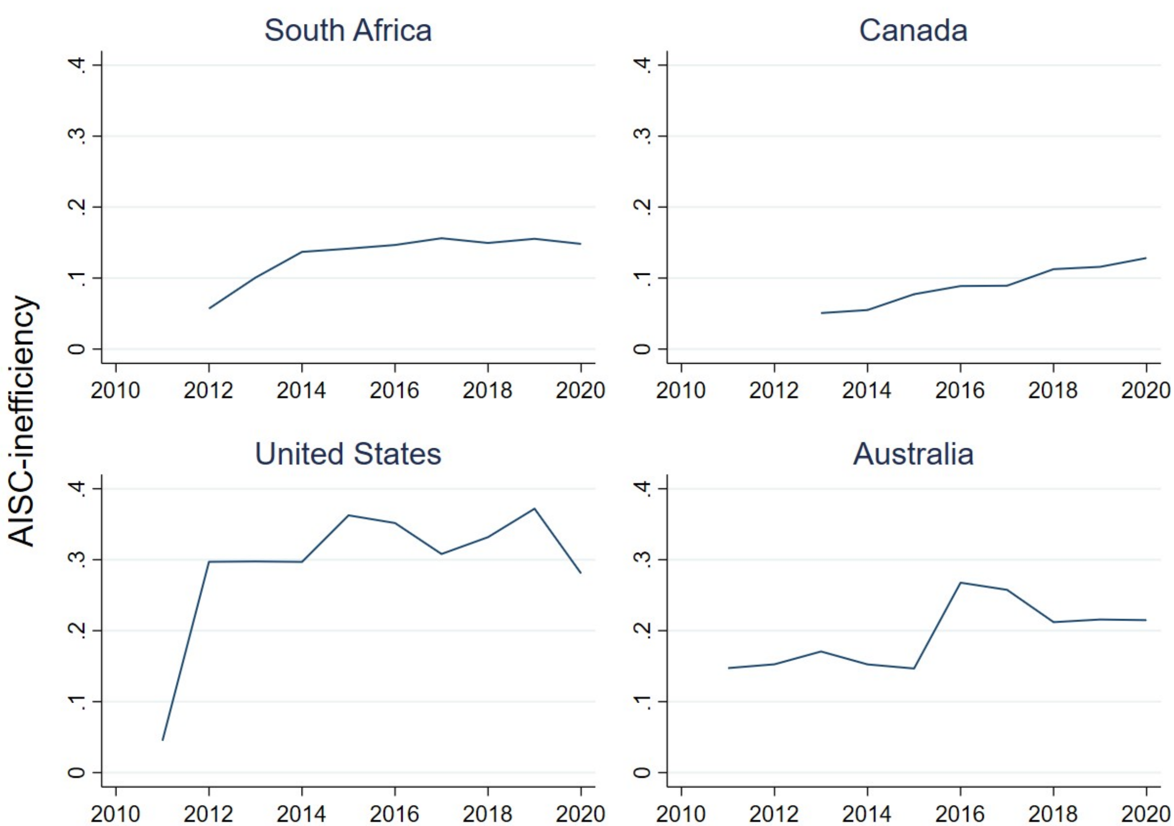

Notes: The four countries are where the largest number of gold mines are located, based on the estimation sample. Collectively, these four countries account for 47 percent of observations in the estimation sample.

TABLE 4 - OLS results: Female board representation and mine-level inefficiency

	(1)	(2)	(3)	(4)
Panel A - AISC-inefficiency				
Share female directors	-0.047	-0.061^{**}	-0.069^{***}	-0.070^{*}
	(0.071)	(0.029)	(0.025)	(0.037)
Adjusted R-squared	0.347	0.513	0.594	0.632
Panel B - C1-inefficiency				
Share female directors	-0.023	-0.043^{*}	-0.044^{*}	-0.070^{**}
	(0.109)	(0.024)	(0.023)	(0.032)
Adjusted R-squared	0.086	0.650	0.652	0.653
Panel C - SC-inefficiency				
Share female directors	-0.126	-0.166^{*}	-0.118^{*}	-0.099
	(0.201)	(0.096)	(0.066)	(0.081)
Adjusted R-squared	0.142	0.520	0.647	0.609
Mine country FE (ω_c)	YES	YES	YES	YES
Year effects (τ_t)	YES	YES	YES	YES
Company FE (σ_j)	No	YES	YES	YES
Company-level controls (X_{jt})	No	No	YES	No
Lagged ROA & asset turnover	No	No	No	YES
Companies (clusters)	41	41	36	37
Mines	90	90	81	81
Observations	352	352	309	297

Notes: OLS regressions with robust standard errors clustered at the company level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The dependent variable is indicated in the panel heading. In all regressions, the observations are weighted by the ownership interest (%) held by company j in mine i at time t. The company-level controls (column 3) are: the log of assets (measured in the company's local currency), and the debt to equity ratio.

FIGURE 4 - Female board representation and AISC-inefficiency (based on Table 4)

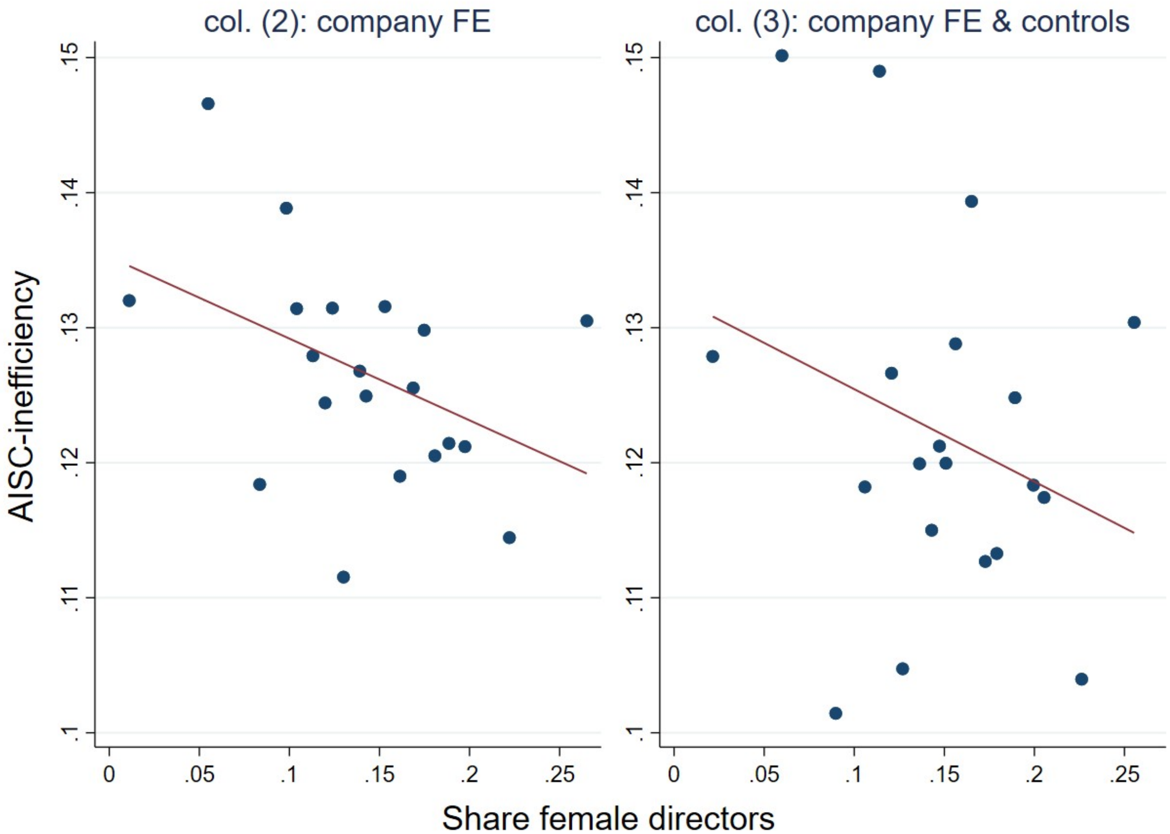

Notes: The plot depicts the conditional relationship between the share of female directors sitting on the parent company's board and the AISC-inefficiency of the mines controlled by the parent company, controlling for mine-country FE, year effects, company FE, and peers' ROA (left panel), as well as two company-level controls (right-panel). The plot groups the observations used in the regressions (Table 4, columns 2 and 3, Panel A) into 20 equal-sized bins. It then computes the mean of (residualised) AISC-inefficiency and (residualised) share female directors within each bin, and displays the 20 data points resulting from this process.

TABLE 5 - IV estimation: Female board representation and mine-level inefficiency

	OLS (1)	2SLS (2)	2SLS (3)	2SLS (4)	2SLS (5)
Panel A - AISC-inefficiency					
Share female directors	-0.094^{***}	-0.154^{***}	-0.145^{***}	-0.155^{***}	-0.185^{**}
	(0.031)	(0.034)	(0.039)	(0.044)	(0.079)
Panel B - C1-inefficiency					
Share female directors	-0.066^{**}	-0.105^{***}	-0.100^{**}	-0.107^{***}	-0.147^{**}
	(0.026)	(0.032)	(0.036)	(0.032)	(0.067)
Panel C - SC-inefficiency					
Share female directors	-0.167^{*}	-0.273^{**}	-0.253^{**}	-0.365^{**}	-0.269^{**}
	(0.093)	(0.123)	(0.114)	(0.164)	(0.135)
Panel D - First-stage equation					
Average share female directors in same:		City	City	Province	Country
		0.850^{***}	0.851^{***}	0.814^{***}	0.984^{***}
		(0.076)	(0.084)	(0.140)	(0.199)
F-statistic (instrument)		124.2	102.5	33.7	24.5
Partial R-squared (instrument)		0.53	0.54	0.32	0.28
Company-level controls	No	No	YES	No	No
Observations	291	291	279	318	328

Notes: 2SLS regressions with robust standard errors clustered at the company level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The dependent variable is indicated in the panel headings. All models include mine country (ω_c), year (τ_t) and company fixed effects (σ_j), as well as peers' ROA. The company-level controls (column 3) are: log of assets, debt-to-equity ratio. In all regressions, the observations are weighted by the ownership interest (%) held by company j in mine i at time t. The instrument used in the 2SLS procedure is a (jackknifed) average share of female directors in the companies located in the same city/province/country as the own company. The F-statistic is the Kleibergen-Paap rk Wald F statistic. The partial R-squared measures the fraction of the total variance in the endogenous regressor (share female directors) that is explained by the instrument, after partialling out the effects of the other exogenous regressors.

TABLE 6 - IV estimation: Additional board-level controls

Additional control(s):	Board size	Share independent directors	Average directors' age	Share new directors	All
	(1)	(2)	(3)	(4)	(5)
Panel A - AISC-inefficiency					
Share female directors	-0.195^{***}	-0.150^{***}	-0.108^{**}	-0.154^{***}	-0.108^{*}
	(0.045)	(0.048)	(0.046)	(0.032)	(0.065)
Panel B - C1-inefficiency					
Share female directors	-0.175^{***}	-0.122^{***}	-0.078^{*}	-0.103^{***}	-0.117^{*}
	(0.045)	(0.045)	(0.047)	(0.031)	(0.062)
Panel C - SC-inefficiency					
Share female directors	-0.227^{*}	-0.262^{*}	-0.178^{*}	-0.283^{**}	-0.157
	(0.136)	(0.141)	(0.104)	(0.122)	(0.138)
Observations	291	291	256	291	256

Notes: 2SLS regressions with robust standard errors clustered at the company level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The dependent variable is indicated in the panel heading. In all regressions, the observations are weighted by the ownership interest (%) held by company j in mine i at time t. All models include mine country (ω_c), year (τ_t) and company fixed effects (σ_j). They also control for the (jackknifed) average ROA in the companies located in the same city as the own company. The instrument used in the 2SLS procedure is a (jackknifed) average share of female directors in the companies located in the same city as the own company.

Table 7 - Female board representation and firm profitability

Dependent variable:	Gross income/assets	Operating income/assets	Pretax income/assets	Net income/assets (=ROA)
	(1)	(2)	(3)	(4)
Dep. var. mean (s.d.)	0.093 (0.101)	0.038 (0.091)	-0.026 (0.230)	-0.044 (0.215)
Share female directors	0.017	0.013	0.204	0.094
	(0.112)	(0.113)	(0.443)	(0.439)
Observations	89	89	89	89

Notes: 2SLS regressions with robust standard errors clustered at the company level. All models include time effects, company FE and the jack-knife average of ROA amongst the firm's city-level peers. Gross income = Sales revenues - COGS. Operating income = Gross income - Operating expenses. Pretax income = Operating income - debt servicing payments. Net income = Pretax income - taxes. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Appendix A - Instrumental variable: First-stage relationships

Dependent variable:	Av. share of fem. dir. in same city ($\overline{Fem}_{jt}^{(C)}$)	(s.e.)	Companies (clusters)	Observations	Within R^2
Share female directors (Fem_{jt})	0.765***	(0.067)	22	82	0.43
Av. ROA in same city	-0.200	(0.148)	19	72	0.01
Log of assets	-0.195	(0.487)	19	72	0.00
Debt to equity ratio	0.688	(2.084)	19	71	0.00
ROA	0.086	(0.475)	19	71	0.00
Asset turnover	-0.016	(0.173)	19	72	0.00
Share new directors	0.067	(0.198)	22	82	0.00
Board size	31.63***	(10.51)	22	82	0.14
Board age	5.085	(8.494)	15	61	0.01
Share independent directors	0.694**	(0.248)	22	82	0.11

Notes: OLS regressions at the company-year level with cluster-robust standard errors in parenthesis. All regressions control for company and year FE. Each row corresponds to a different regression. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Appendix B - Alternative stochastic frontier models

	Alternative inefficiency term	Specification of cost function		
		No controls	Linear trends	Quadratic term in output
	(1)	(2)	(3)	(4)
Panel A - AISC-inefficiency				
(1a) Cost function (DV: ln AISC):				
ln(Gold output)	0.824***	0.916***	0.820***	1.058***
	(0.039)	(0.043)	(0.041)	(0.184)
ln(Gold output)^2				-0.026
				(0.021)
Observations	916	986	916	916
(1b) Inefficiency term (u_{irt}):				
Mean (s.d.)	0.122 (0.123)	0.186 (0.205)	0.139 (0.168)	0.149 (0.184)
Correlation with baseline u_{irt}	0.98	0.84	0.99	0.99
(2) 2SLS regression (DV: u_{irt}):				
Share female directors	−0.115***	−0.157***	−0.139***	−0.152***
	(0.025)	(0.040)	(0.032)	(0.040)
Panel B - C1-inefficiency				
(1a) Cost function (DV: ln C1):				
ln(Gold output)	0.800***	0.837***	0.792***	1.277***
	(0.041)	(0.031)	(0.046)	(0.176)
ln(Gold output)^2				-0.055**
				(0.021)
Observations	1251	1351	1251	1251
(1b) Inefficiency term (u_{irt}):				
Mean (s.d.)	0.113 (0.111)	0.165 (0.176)	0.202 (0.235)	0.150 (0.181)
Correlation with baseline u_{irt}	0.99	0.82	0.74	0.94
(2) 2SLS regression (DV: u_{irt}):				
Share female directors	−0.084***	−0.123***	−0.130***	−0.122**
	(0.024)	(0.036)	(0.029)	(0.052)
Panel C - SC-inefficiency				
(1a) Cost function (DV: ln SC):				
ln(Gold output)	0.990***	1.172***	1.028***	1.498***
	(0.069)	(0.074)	(0.079)	(0.356)
ln(Gold output)^2				-0.054
				(0.037)
Observations	557	588	557	557
(1b) Inefficiency term (u_{irt}):				
Mean (s.d.)	0.145 (0.127)	0.387 (0.521)	0.101 (0.112)	0.191 (0.220)
Correlation with baseline u_{irt}	0.97	0.49	0.85	0.98
(2) 2SLS regression (DV: u_{irt}):				
Share female directors	−0.156**	−0.156	−0.175**	−0.256**
	(0.072)	(0.103)	(0.087)	(0.115)

Notes: The sections (1) report ML regressions with robust standard errors clustered at the mine level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. All regressions control for province-level quadratic trends and include the full set of mine-level controls (unless otherwise stated). The sections (2) report the mean (standard deviation) of the inefficiency terms derived from the stochastic frontier analyses reported in the sections (1). They also report the correlation coefficients with the baseline inefficiency terms obtained from the models shown in Table 2 (and used in the main analysis). The sections (3) report the estimated coefficient on *Fem* obtained from 2SLS regressions as in Table 5, column 2, using the alternative estimates of the inefficiency term as dependent variables. In column (1), the inefficiency term is estimated using the method of Battese and Coelli (1988) instead of the method of Jondrow et al. (1982). In column (2), the technology controls (X_i, see Table 2) are omitted from the cost function. In column (3), the cost function includes province-level linear (instead of quadratic) trends ($\rho_r + \theta_r t$). In column (4), the cost function includes a quadratic term in output.

Appendix C - Alternative measures of female boardroom representation

Dependent variable:	AISC-inefficiency (1)	C1-inefficiency (2)	SC-inefficiency (3)
Panel A: At least one female director (dummy)			
$I(Fem \geq 1)$	-0.124^{*}	-0.084	-0.219^{**}
	(0.069)	(0.055)	(0.094)
Panel B: More than one female director (dummy)			
$I(Fem > 1)$	-0.133^{*}	-0.091^{*}	-0.236^{*}
	(0.068)	(0.052)	(0.147)

Notes: 2SLS regressions with robust standard errors clustered at the company level, as in Table 5, column 2. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The number of observations is always 291.

Appendix D - Alternative specifications of equation (3)

Specification:	No ω_c (1)	Mine continent FE instead of ω_c (2)	Company-city FE instead of σ_j (3)	Company-level linear trends (4)
Panel A - AISC-inefficiency				
Share female directors	-0.185^{*}	-0.235^{***}	-0.119^{***}	-0.112^{**}
	(0.095)	(0.075)	(0.035)	(0.046)
Panel B - C1-inefficiency				
Share female directors	-0.155^{***}	-0.211^{***}	-0.119^{***}	-0.132^{**}
	(0.024)	(0.050)	(0.046)	(0.067)
Panel C - SC-inefficiency				
Share female directors	-0.217	-0.241	-0.232^{**}	-0.067
	(0.188)	(0.153)	(0.106)	(0.043)

Notes: 2SLS regressions with robust standard errors clustered at the company level, as in Table 5, column 2. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. In column (1), the mine-country fixed effects (ω_c) are omitted from the regression. In column (2), the mine-country fixed effects are replaced with mine-continent fixed effects. In column (3), the company fixed effects (σ_j) are replaced with company-city fixed effects. Column (4) adds company-level linear trends ($\sigma_j t$). The number of observations is always 291

Appendix E - Alternative inference procedures and clustering structures

Clustering level	Statistic	Dependent variable: AISC-inefficiency	C1-inefficiency	SC-inefficiency
	Parameter estimate (2SLS)	-0.154	-0.105	-0.273
Company	z-statistic	-4.58	-3.31	-2.21
	p-value, normal	(0.000)	(0.001)	(0.027)
	p-value, WRE bootstrap	(0.010)	(0.012)	(0.260)
	z-statistic, WRE bootstrap-c	-2.32	-2.19	-1.67
	p-value	(0.005)	(0.009)	(0.125)
Company & Mine country	z-statistic	-2.54	-2.50	-1.69
	p-value, normal	(0.011)	(0.013)	(0.092)
	p-value, WRE bootstrap	(0.120)	(0.189)	(0.490)
	z-statistic, WRE bootstrap-c	-1.76	-1.57	-1.47
	p-value	(0.016)	(0.096)	(0.199)
Company & Year	z-statistic	-5.08	-4.52	-2.45
	p-value, normal	(0.000)	(0.000)	(0.014)
	p-value, WRE bootstrap	(0.042)	(0.065)	(0.149)
	z-statistic, WRE bootstrap-c	-1.89	-1.83	-1.51
	p-value	(0.047)	(0.067)	(0.071)

Notes: There are 41 company clusters and 291 observations. The bootstrap dimension for two-way clustering is always the coarser one (mine country and year, respectively). The WRE bootstrap procedure always employs 9,999 replications using the Rademacher distribution (Webb distribution for company & year clustering). Equal-tail *p*-values are reported in parenthesis.