THESIS / THÈSE

DOCTOR OF ECONOMICS AND BUSINESS MANAGEMENT

Essays on Mechanism Design and Family Economics

Woitrin, Francois

Award date:
2023

Awarding institution:
University of Namur

Université de Namur

Thèse pour l'obtention du titre de Docteur en Sciences Economiques

Essays on Mechanism Design and Family Economics

Thèse défendue publiquement le 13 septembre 2023 par

François Woitrin

Supervisors

Prof. Jean-Marie Baland (Université de Namur)
Prof. Guilhem Cassan (Université de Namur)

Jury Members

Prof. Siwan Anderson (University of British Columbia)
Prof. Paula Gobbi (Université Libre de Bruxelles)
Prof. Pauline Rossi (Ecole Polytechnique - CREST)

President

Prof. Lorenzo Trimarchi (Université de Namur)

Remerciements

Je tiens tout d'abord à exprimer ma profonde gratitude envers mes superviseurs, Jean-Marie Baland et Guilhem Cassan, pour leur attention constante, leurs conseils éclairés, et la confiance qu'ils ont placée en moi tout au long de ce passionnant parcours académique.

Je souhaite également exprimer ma reconnaissance envers Benoît Decerf, qui a été non seulement mon superviseur pendant une période de ma thèse, mais aussi le co-auteur et instigateur du premier chapitre. Dès les premières semaines de cette aventure, sa confiance en mon travail a été un moteur précieux.

Les membres du jury, Siwan Anderson, Paula Gobbi et Pauline Rossi, méritent toute ma gratitude pour leur temps consacré à la lecture attentive de ma thèse, ainsi que pour leurs commentaires et contributions qui ont grandement enrichi mon travail. En outre, je tiens à exprimer ma reconnaissance envers François Libois, Nicolas Debarsy, ainsi que tous les autres académiques du département, pour leurs aides et suggestions avisées au fil de ces années.

Je tiens à remercier chaleureusement toutes les personnes avec lesquelles j'ai eu la chance de collaborer dans le domaine de l'enseignement. Leurs visions et expériences m'ont apporté une perspective unique. Je réserve une mention spéciale à Jean-Marie Cheffert et Benoît Decerf pour l'opportunité extraordinaire qu'ils m'ont offerte en me permettant de suppléer le cours de Stratégie et Décisions Économiques pendant deux ans. Cette reconnaissance s'étend également à Stéphanie Weynants, Vincenzo Verardi, Catherine Guirkinger, Wouter Gelade, Anna Kiriliouk, Germain Van Bever, Nathan Uyttendaele, Patrick Foissac, Christelle Hoorelbeke, Sophie Pondeville, Hélène Laurent, Jolan Mohimont, Modest Dayé, Ludovic Bequet, Doux Baraka, Pierre Laurent, Luca Faré, Patricia Nisole, Eli Agba, Nicolas van Keirsbite, Louise Schraverus, Coralie Franc et à Fabrice Orban pour son assistance lors de mes défis informatiques.

Je n'oublie pas ceux qui ont contribué à rendre mon temps à la faculté mémorable. Mes sincères remerciements vont à Ludovic et Marine, mes collègues de bureau, ainsi qu'à Joey, Camille, Lhorie, Jérôme, Henry, Olivier, Jolan, Modeste, Christian, Nathalie, Marie, Baptiste, Pierre, Clarice, Coline, Marie, Quentin, Perrin, Hugues, Andréa, Pablo, Anna, Giorgio, Etienne, Nitin, Jolan, Marcus et Nicolas pour les nombreux moments d'échange et de camaraderie partagés autour d'une tasse de café ou d'une bière.

Enfin, je voudrais exprimer toute ma reconnaissance envers ma famille. Maman, Papa, Geoffroy, Isabelle, Pierre-Henri, Marine, Juliette et Bonne-maman, votre soutien indéfectible au cours de ces six années a été essentiel à ma réussite académique.

Pour finir, mes amis, et particulièrement l'ASBLP, méritent une mention spéciale pour les moments inoubliables que nous avons partagés en dehors des heures de travail. Votre amitié a apporté une précieuse équilibre à ma vie pendant cette période intense de recherche.

Contents

List of Figures

List of Tables

Introduction

Mechanism design stands as a fundamental concept in economics, addressing the complex challenge of making collective choices based on the preferences of a group of individuals. It has practical applications in diverse scenarios, ranging from the allocation of seats in public schools, to the distribution of goods through auctions or the design of voting and political systems. Central to mechanism design is the arrangement of incentives that guide individuals to strategically interact and select outcomes aligned with desirable properties, which are formalized as "social choice correspondences".

Consider, for instance, the allocation of spots in public schools. A mechanism is needed to determine which students are assigned to which schools, taking into account their preferences and the available seats. In this context, a well-designed mechanism could ensure that students reveal their true preferences, allowing the system to allocate seats efficiently, while satisfying a fairness criteria for instance. However, preferences are often private information, hidden from the decision-makers. In the previous example, this means that a decision-maker can not know the exact preferences of every student about every existing school. This private information complicates the process of determining which students should be placed where, based on their genuine preferences. Instead of directly assessing whether equilibrium outcomes satisfy desirable properties (which is what game theory focuses on), the focus shifts to constructing a mechanism that induces individuals to choose the desired outcome. A successful mechanism has the ability to incentivize strategic behavior that leads to outcomes aligning with the desired properties. Individual incentives, collective welfare, and strategic interactions forms the heart of mechanism design.

A central goal of mechanism design therefore lies in the comparison of different mechanisms to discern their effectiveness in achieving desired outcomes. In order to do this, the first step must be to anticipate the strategies individuals might adopt when facing each mechanism based on their preferences. These predictions are made from concepts shared with game theory like Nash equilibrium or undominated strategies, for instance. One must then find the outcomes generated by each mechanism when agents play the predicted strategies, before assessing if those outcomes meet the desirable properties. In real-world scenarios, solution concepts might offer multiple predictions for the strategies individuals could adopt within a mechanism. Consequently, some of these predictions may yield outcomes satisfying the desired property, while others might not. This diversity of predictions raises questions about how mechanisms should be compared. For instance, when considering different mechanisms, how does one assess the performance of a mechanism producing one desirable and one undesirable outcome compared to another generating four desirable and two undesirable outcomes?

In this first chapter, we focus on assessing and comparing mechanisms based on their ability to satisfy specific properties, such as Pareto efficiency or stability, particularly when multiple equilibria are at play. Through a process of axiomatization, we characterize three distinct criteria for comparing mechanisms. The first criterion is intuitive and based on highly compelling axioms but is also incomplete (in the sense that it does not allow to compare many mechanisms) and difficult to work with. In acknowledging its limitations, we develop two additional criteria, building on the foundations of the first, while introducing robustness axioms that increase their applicability and extend their scope. These new criteria offer a more comprehensive perspective on the comparison of mechanisms, giving us a deeper understanding of their strengths and limitations across various preference profiles. In summary, our research engages in an exploration of mechanism design's theoretical and practical aspects, offering a perspective on the evaluation and comparison of mechanisms across various settings.

Moving from the realm of mechanism design, we shift to another important societal issue that significantly impacts gender equality, family dynamics, and societal development: the prevalence of son preference and its associated "stopping rule."

In many developing regions, a preference for sons over daughters is a deeply rooted element of the social structures of societies. This phenomenon goes beyond being a mere cultural tradition exerting a substantial impact on communities, as it significantly influence reproductive patterns, family structures, or gender role expectations. Marriages for instance are often concluded with the implicit goal of bearing sons who will ensure the continuation of family legacies and traditions; in some communities, the birth of a son is grandly celebrated for it symbolizes the continuation of the family name and heritage. Inheritance patterns further compound these disparities as, in many cultures, sons are given a larger share of family assets, reinforcing the notion of their pivotal role in perpetuating the family lineage. They are also perceived as future breadwinners and supporters for their parents in old age.

The consequences of son preference are manifold and may vary across regions, but often extend far beyond only shaping societal norms. At the individual level, daughters frequently find themselves at a disadvantage in terms of access to essential resources. For instance, limited health care resources may be allocated to sons over daughters, affecting overall well-being. Additionally, unequal access to education can perpetuate gender-based disparities in skills and productivity, contributing to cycles of poverty, restricting women's activities and reducing their opportunities for personal and economic advancement. Overall, the consequences of this preference extend throughout society, leading to persistent gender disparities and presenting substantial challenges to achieving development objectives.

Among the various consequences of son preference, this thesis focuses on what is known in the literature as the "stopping rule" a central behavior when discussing son-favoring fertility practices. The stopping rule refers to the deliberate decision by parents to continue having pregnancies until they achieve a specific number of sons, regardless of the total desired family size. This behavior manifests itself into two distinct yet interconnected behaviors: instrumental births and sex-selective abortion. Instrumental births encompass situations where parents continue to have children until reaching the desired number of boys. This can result in larger-than-planned families and, at times, lead to unintended consequences for the well-being of both the parents and the children (overcrowded living conditions, child labor to contribute to the family's income, or maternal health for instance). On the other hand, sex-selective abortion involves the termination of pregnancies based on the gender of the fetus. When parents desire male children, they may opt for sex-selective abortion if prenatal testing reveals a female fetus. The consequences of the stopping rule and sex-selective abortion reverberate through multiple layers of society. At the individual level, these practices can perpetuate gender inequalities and have significant repercussions for the well-being of both sons and daughters. Sons may bear the burden of fulfilling societal expectations, while daughters may face diminished access to resources and opportunities, leading to cycles of disadvantage. Within households, the pursuit of male offspring can influence family dynamics, inheritance patterns, and intra-household decision-making processes. Sons may receive preferential treatment in terms of education, health care, and other opportunities, while daughters might experience limited agency and reduced access to essential resources. On a broader societal scale, sex-selective abortion contribute to skewed sex ratios, which can lead to imbalances in the marriage market and social stability. The resulting gender disparities can hinder economic growth and development by under-utilizing the potential of women and perpetuating a cycle of inequality.

By shedding light on the multifaceted impacts of these practices, this thesis aims to contribute to a deeper understanding of the challenges posed by son preference and pave the way for informed policy interventions that promote gender equity, individual well-being, and sustainable societal development. The second chapter introduces the theoretical underpinnings of the stopping rule, emphasizing its twofold manifestation instrumental births and sex-selective abortion. By viewing these practices within an integrated framework, this chapter sets the stage for a comprehensive understanding of this complex phenomenon. We propose measures of detection and offer a novel perspective for gauging the prevalence of these behaviors across diverse countries. The third chapter delves deeper into the repercussions of the stopping rule. It highlights how this behavior generates sibling competition, resulting in gender-specific disparities in mortality rates. This empirical investigation quantifies the extent to which the stopping rule contributes to girls' mortality, uncovering surprising insights into its impact on the less privileged segments of the population of some countries.

The **first chapter**, *Criteria to compare mechanisms that partially satisfy a property: an axiomatic study*, co-authored with Benoît Decerf, introduces and axiomatize three criteria to be used as tools for making informed choices between mechanisms. In the illustrative example,we shift to the comparison of two matching mechanisms for school seat allocation: DA^2 and BOS^2, constrained versions of the Deferred Acceptance

(DA) and Boston (BOS) mechanisms. These mechanisms are compared based on their stability with respect to the "stable" social choice correspondence. The comparison highlights the challenges of our first criterion, the Proportion criterion, in assessing stability. Even in this simplified setting, determining the proportion of stable equilibria for each mechanism involves complex calculations. The study then introduces the PHO and the PHO^* criteria (PHO for Profiles of Homogeneous Outcomes), which overcome the limitations of the Proportion criterion. The increased discriminatory power of PHO allows for a clearer comparison of stability between DA^2 and BOS^2. The example is used to show that DA^2 performs better than BOS^2 in terms of stability according to the PHO criterion. The paper also establishes relationships between PHO and PHO^*, enabling deductions about the relative performance of mechanisms based on these criteria.

The **second chapter**, *Sex-Selective Abortions and Instrumental Births as the two faces of the Stopping Rule. New measures and world evidence*, co-authored with Jean-Marie Baland and Guilhem Cassan, highlights the prevalence of the stopping rule in a large set of countries. This includes countries widely known for their strong preference for sons such as India, Bangladesh and Nepal, but also countries which are much less often mentioned in the literature like Armenia, Albania, Azerbaijan or Ukraine. These countries exhibit a wide range of gender preference biases, with desired sex ratios varying between 109 and 232, always significantly surpassing 'natural' sex ratios. We also show that instrumental births overwhelmingly contribute to the prevalence of the stopping rule, accounting for more than two-thirds of its impact. In stark contrast to previous studies that predominantly focus on sex-selective abortion, our comprehensive approach unveils a far more extensive influence of instrumental births. As expected, girls are consistently more likely to be instrumental. For instance, in Armenia, 64.5% of girls are considered instrumental compared to 25.3% of boys, and in India, the corresponding figures are 53.4% and 34.2%. Finally, a more detailed analysis within India reveals distinct patterns across states and social groups. Gender-biased preferences follow well-established caste hierarchies, with higher castes exhibiting stronger biases. For instance, among high castes, the desired sex ratio reaches an average of 246%, while among Muslims, it falls to 140%.

The **third chapter**, titled *Stopping Rule and Girls Mortality: Insights from South Asian Nations*, delves into the implications of the fertility consequences associated with the stopping rule, as explored in the preceding chapter, on the under-five mortality rate. Within this context, we define the mechanism involving the indirect impacts of the gender-biased stopping rule on health outcomes as 'passive discrimination'. The investigation reveals compelling evidence of substantial gender-specific differentials in the levels of intra-household competition experienced by boys and girls. Moreover, these disparities in competition intensity exhibit variations upon household characteristics, notably wealth. Concurrently, the study establishes a statistically significant influence of sibling competition on mortality rates. For instance, the findings show that in India, the presence of an additional sibling during the initial five years of a childs life increase the likelihood of her death before the age of five by an average of approximately 12.61 percentage points. This effect is particularly pronounced within the poorest strata of the population. From these two estimations, we quantify the association between passive discrimination and female mortality. Over the period spanning from 1980 to 2015, the analysis attributes over 2,500,000 girl deaths in India directly to the passive discrimination mechanism. Analogously, the estimated figures for Bangladesh, Pakistan, and Nepal amount to 120,000, 90,000, and 25,000 respectively. These represent a substantial proportion of girls' deaths, with India facing an impact that accounts for up to 20% of under-five female mortality. As expected, most of these cases occur within the poorest segments of the population.

Chapter 1

Criteria to compare mechanisms that partially satisfy a property: an axiomatic study

Joint with Benoît Decerf

The paper was published in the journal Social Choice and Welfare (2022)

Abstract: We study criteria that compare mechanisms according to a property (e.g., Pareto efficiency or stability) in the presence of multiple equilibria. The multiplicity of equilibria complicates such comparisons when some equilibria satisfy the property while others do not. We axiomatically characterize three criteria. The first criterion is intuitive and based on highly compelling axioms, but is also very incomplete and not very workable. The other two criteria extend the comparisons made by the first and are more workable. Our results reveal the additional robustness axiom characterizing each of these two criteria.

1.1 Introduction

From the assignment of seats at public schools to the allocation of goods against payment in auctions, economics repeatedly faces the problem of choosing among outcomes based on the preferences of a set of agents over these outcomes. To guide such collective choices, outcomes are often sorted according to desirable properties, formalized as *social choice correspondences*. If the agents' preferences are known, the set of outcomes can, for example, be sorted into subsets of Pareto efficient and Pareto inefficient outcomes, or in some applications, into subsets of "fair" and "unfair" outcomes.

Of course, preferences are often private information which makes it impossible for the social planner to directly compute whether outcomes satisfy a desirable property. Instead, the planner must setup a mechanism through which agents interacting strategically determine the selected alternative. Guiding the planner's design requires determining which mechanism better provides agents with the incentives to select strategies that, given their preferences, lead to the selection of outcomes satisfying the desirable property.

The comparison of competing mechanisms then follows a three-step procedure. The first step consists in predicting the strategies agents might use in each mechanism as a function of their preferences. Formally, these predictions are captured by *solution concepts* such as undominated strategy, Nash equilibrium, or dominant strategy. Second, one must compute the outcomes selected by the mechanisms when agents play the strategies predicted by the solution concepts selected in the first step. Finally, one must evaluate the resulting outcomes according to the property of interest. Example applications of this three-step procedure can be found in Ergin and Sönmez (2006) and Abdulkadiroğlu et al. (2011).

In practice, given a preference profile, it is common for solutions concepts to make multiple predictions about the strategies agents might use in a mechanism. When this is the case, some of the predicted outcomes might satisfy the property of interest, whereas others might not. When this is the case, it is often unclear how mechanisms should be compared. For example, on a given preference profile, how does a mechanism with one desirable and one undesirable outcomes compare with a mechanism with four desirable and two undesirable outcomes?

In this paper, we propose, characterize and compare three criteria to perform such comparisons. First, the "Proportion" criterion compares, on a profile-by-profile basis, the *fraction* of desirable outcomes reached by each mechanism (the higher the fraction, the better the mechanism performs in terms of the property at stake). We show that this natural criterion is characterized by three compelling axioms. Unfortunately, the Proportion criterion only provides a very partial ranking of mechanisms and often concludes that mechanisms cannot be compared. Moreover, this criterion is not very workable because it requires counting the number of equilibria and identifying the fraction of desirable equilibria. Doing so becomes increasingly difficult as the number of equilibria grows.

Our two other criteria improve on both limitations and therefore constitute our main contribution. Both criteria satisfy the same three axioms as the Proportion criterion, and therefore agree with it on all pairs that the Proportion criterion is able to rank. To provide more complete orderings, each of these two additional criteria also satisfy an additional robustness axiom. Loosely put, these two mirror robustness axioms require that a comparison between two mechanisms would not be altered if both mechanisms had one additional desirable (undesirable) outcome.

Importantly, the two "extended" criteria compare mechanisms by focusing on preference profiles for which outcomes are either all desirable, or all undesirable. By doing so, they yield more affirmative comparisons because they are not necessarily bogged down by the existence of a few preference profiles for which the proportions of desirable outcomes are reversed. Moreover, these criteria do not require counting the number of equilibria nor computing the fraction of desirable outcomes.

Of course, the strength of an affirmative comparison between two mechanisms depends on the criterion used. One can be more confident that a mechanism will perform better than another when they can be ranked by the Proportion criterion than when this can only be done using our two other criteria. Yet, when the Proportion criterion is silent, comparisons in terms of our dual criteria provide interesting indications about the respective performance that should be expected from two alternative mechanisms. In other words, the improvement on the limitations associated with the Proportion criterion comes at some cost. Our robustness axiom can lead to comparisons between mechanisms that are somewhat more debatable. This axiom can be viewed as capturing the cost of increasing the completeness of the partial order.

We illustrate the different discriminative powers and workabilities of these criteria for the comparison of the stability of two school choice mechanisms on a narrow domain.

The paper is organized as follows. We integrate our work in the literature in Section 1.2. We present the framework in Section 1.3. We derive axiomatically our criteria and discuss their shortcomings in Section 1.4. We then illustrate how those criteria can be use in the school choice problem in Section 1.5 and conclude in Section 1.6.

1.2 Related Literature

Our three criteria use a "profile by profile" approach to compare mechanisms, which is common in the literature on voting procedures (Dasgupta and Maskin (2008); Gerber and Barberà (2016); Arribillaga and Massó (2015)).[1] This approach is also common in the matching literature. Whereas our paper focuses on comparing the properties of outcomes, the matching literature has proposed a number of criteria to compare the manipulability of mechanisms. Pathak and Sönmez (2013) for example rank mechanisms by comparing the set of preference profiles for which the mechanisms admit a truthful Nash equilibirum. If a mechanism admits a truthful Nash equilibrium in every profile for which another mechanism also does, then Pathak and Sönmez (2013) conclude that the latter is less manipulable than the former. Similarly, Andersson et al. (2014) study manipulability by comparing the *number* of preference profiles at which each mechanism is manipulable. This type of manipulability comparisons avoids the issue induced by multiple solutions since it relies on binary evaluations: For any

[1] In Gerber and Barberà (2016), the solution concept is "iterated elimination of weakly dominated strategies" and the correspondence is the possibility of agenda manipulation. In Dasgupta and Maskin (2008), the solution concept is "truthful revelation" and the correspondence is a collection of five voting properties.

given preference profile, either a mechanism is manipulable or it is not.

In contrast, the multiplicity issue is key when evaluating the efficiency or fairness of outcomes. For example, Chen and Kesten (2017) compare school choice mechanisms with respect to the stability of their Nash equilibrium outcomes. The criterion implied by their analysis relies on the comparison of the number of stable equilibria in each type profile. Ergin and Sönmez (2006) show that the multiple equilibria of the Boston mechanism are all Pareto dominated by that of the Deferred Acceptance mechanism.

When mechanisms do not perfectly satisfy a property of interest, another approach consists in comparing them using a criterion formalizing “by how much" each solution violates the property. In the case of stability, which requires the absence of *blocking pairs*, Combe et al. (2017), Abdulkadiroglu et al. (2019), Dogan and Ehlers (2020b) and Bonkoungou and Nesterov (2020) compare mechanisms by measuring, in different ways, the number of blocking pairs or the number of players participating to a blocking pair in each profile. Dogan and Ehlers (2020a) axiomatically characterize criteria for stability comparisons based on axioms specific to this property. Current research along this approach has so far abstracted from the multiplicity issue that we aim at tackling here.

Going even further away from the binary nature of social choice correspondences, which only acknowledge two desirable or undesirable categories, some authors propose to compare mechanisms using fine-grained normative tool, e.g., a social welfare function. Fleurbaey (2012) axiomatically characterize a criterion that compares how alternative mechanisms perform in the light of a fined-grained ranking of outcomes. Again, the setting considered by that author abstracts from the multiplicity issue.

Finally, our work shares important similarities with the literature on the measurement of predictive success (Selten (1991)). We derive criteria that compare mechanisms as a function of their ability to yield outcomes that are selected by a correspondence. That literature derives rules that compare theories as a function of their ability to make predictions that are in line with observations. There are fundamental differences between these two objectives, which imply that our criterion are unrelated to these rules. Indeed these differences in objectives makes the relevant primitives different as well.[2]

1.3 Framework and notation

This section introduces the terminology and notation for our axiomatic results. We let $N = \{1, \ldots, n\}$ denote the set of players, and $o \in O$ denote the set of outcome. Each player $i \in N$ is characterized by a type $y_i \in Y_i$, e.g., the player's

[2]The relevant primitives for our criteria are the numbers of equilibria that yield an outcome that is (resp. not) selected by the correspondence. In contrast, the relevant primitives for these rules include the "hit rate", i.e. the fraction of observations predicted by the theory, and the "area", i.e. the fraction of potential outcomes predicted by the theory. Neither the "hit rate" nor the "area" are relevant primitives for our criteria. We provide here the intuition why the "area" is not a relevant primitive for our criteria. The following two mechanisms have different area but should be considered equivalent by our criteria. The first mechanism has a unique equilibrium that yields an outcome that is selected by the correspondence. The second mechanism has multiple equilibria, all of which yield outcomes that are selected by the correspondence.

preference over the outcomes in O. A type profile is denoted by $y \in Y := \times_{i \in N} Y_i$. Let $X\colon Y \to 2^O$ be a **social choice correspondence**, sometimes *correspondence*, for short.

A **mechanism** is a game form $M\colon S \to O$ that associates every strategy profile $s \in S := \times_{i \in N} S_i$ with an outcome in O, where S_i is the *finite* strategy space of $i \in N$. The set of mechanisms is $\mathcal{M}$ ($\mathcal{M}$ includes both direct and indirect mechanisms).

Let $C\colon Y \times \mathcal{M} \to 2^S$ denote a **solution concept**. The set $C(y, M)$ corresponds to the set of strategy profiles that C predicts could be played in mechanism M when the type profile is y. As is common, we henceforth refer to $C(y, M)$ as the set of **equilibria** of M under C when the type profile is y (whether or not C is an "equilibrium" solution concept). Since we assume that strategy spaces are finite, the number of equilibria is always finite. We focus on solution concepts that admit at least one equilibrium for each type profile. The set of such solution concepts is $\mathcal{C}$.

For a given correspondence X, let $\succeq$ be a **partial order** on $\mathcal{M} \times \mathcal{C}$. A partial order is a binary relation that is reflexive, asymmetric, and transitive.[3] The relation $(M, C) \succeq (M', C')$ indicates that mechanism M satisfies the property corresponding to X at least as well as mechanism M' when the former is played according to solution concept C and the latter according to solution concept C'. The symmetric and anti-symmetric relations, i.e., $(M, C) \succ (M', C')$ and $(M, C) \sim (M', C')$, are defined accordingly. Because $\succeq$ is partial, there may exists pairs $[(M, C), (M', C')]$ for which the relation is undefined.

Observe that we require the partial order to compare pairs

[3] In particular, the weak relation is transitive.

(M, C), (M', C') that are potentially based on different solution concepts. This recognizes the fact that the behavior and coordination possibilities of players may depend on the mechanism. This is especially true when one of the mechanisms under consideration admits dominant strategies whereas the other does not (in which case, it is reasonable to use dominant strategies as a solution concept for the mechanisms where the latter is non-empty, and use the next best solution concept for the other mechanism, see, e.g.,Ergin and Sönmez (2006); Abdulkadiroğlu et al. (2011)).

Our objective is to identify partial orders satisfying compelling properties. Throughout, we restrict our attention to partial orders that satisfy an independence property we call *Outcome Neutrality.* This property forces partial orders to compare mechanisms based only on the *number* of equilibria whose outcome are selected (or not) by the social choice correspondence.[4] This captures the idea that the only aspect of equilibrium outcomes that matters to $\succeq$ is whether or not they are selected by the correspondence X. For any set A, we let $\#A$ denote the cardinality of set A.

Axiom 1 (*Outcome Neutrality*).
For all $C, C' \in \mathcal{C}$ and all $M, M' \in \mathcal{M}$, if for all $y \in Y$ we have

(i) $\#\{s \in C'(y, M') \mid M'(s) \in X(y)\} = \#\{s \in C(y, M) \mid M(s) \in X(y)\}$, *and*

(ii) $\#\{s \in C'(y, M') \mid M'(s) \notin X(y)\} = \#\{s \in C(y, M) \mid$

[4] This property assumes that all equilibria count the same. This is a natural assumption if one believes that all equilibria are equally likely to occur.

$M(s) \notin X(y)\}$,

then $(M, C) \sim (M', C')$.

All partial orders satisfying *Outcome Neutrality* can be reformulated as partial orders over particular "*counting*" functions. Any pair (M, C) defines an associated **counting function** F that associates any y with a function $F(y)$ such that $F_0(y) \coloneqq \#\{s \in C(y, M) \mid M(s) \notin X(y)\}$ and $F_1(y) \coloneqq \#\{s \in C(y, M) \mid M(s) \in X(y)\}$. When no confusion on the types profile is possible, we simply write the components of the function F_0 and F_1.

Outcome Neutrality implies that any two (M, C) and (M', C') whose associated functions F and F' are the same perform equally well in terms of correspondence X (formally, $(M, C) \sim (M', C')$ whenever $F = F'$). Therefore, the partial order $\succeq$ on domain $\mathcal{M} \times \mathcal{C}$ is equivalent to a partial order on domain $\mathcal{F} = \{F : Y \to Z\}$, where $Z = \{(z_0, z_1) \in \mathbb{N}_0^2 | z_0 + z_1 \geq 1\}$. Observe that set Z is unbounded, a feature that is necessary for some of our results.[5]

Slightly abusing the notation, we also denote the latter partial order by $\succeq$. For the sake of improved readability, all remaining properties on the partial order are expressed on domain $\mathcal{F}$.

[5]In particular, Parts 2 of Theorems 1, 2 and 3 require the construction of intermediate mechanisms that may have, for some type profiles, more numerous equilibria than the number of equilibria of the mechanisms being compared. However, Parts 1 of Theorems 1, 2 and 3 do not require Z to be unbounded.

1.4 Criteria

We start by presenting three basic axioms for partial orders. When no confusion is possible, we ignore the role of solution concepts and simply say that we compare two mechanisms. Also, we write that an equilibrium is "in X" ("not in X") if its outcome is selected (not selected) by correspondence X. Finally, we say that two mechanisms M and M' are *equivalent* on a type profile y if $F_0(y) = F_0'(y)$ and $F_1(y) = F_1'(y)$.

Our first axiom, *Domination*, requires that if two mechanisms are equivalent on all but one type profile for which all the equilibria of one mechanism are in X whereas all the equilibria of the other mechanism are not in X, then the former performs better than the latter in terms of X.

Axiom 2 (*Domination*)**.** *For all $F, F' \in \mathcal{F}$, if (i) $F_1(y^*) = 0$ and $F_0'(y^*) = 0$ for some y^*, and (ii) $F'(y) = F(y)$ for all $y \neq y^*$, then $F' \succ F$.*

Second, *Monotonicity* captures the idea that a larger number of equilibria not in X does not improve performance, while a larger number of equilibria in X does not worsen it. If two mechanisms are equivalent on all but one type profile for which they are not exactly equivalent because one mechanism has either one more equilibrium in X or one less equilibrium not in X than the other mechanism, then the axiom concludes that the former performs weakly better in terms of X.

Axiom 3 (*Monotonicity*)**.** *For all $F, F' \in \mathcal{F}$, if (i) for some y^* we have either $F_0'(y^*) = F_0(y^*)$ and $F_1'(y^*) = F_1(y^*) + 1$, or $F_0(y^*) = F_0'(y^*) + 1$ and $F_1'(y^*) = F_1(y^*)$, and (ii) $F'(y) = F(y)$ for all $y \neq y^*$, then $F' \succeq F$.*

These two axioms are based on demanding preconditions and therefore, on their own, only impose relatively weak restrictions on $\succeq$. Hence, many implausible partial orders are not ruled out by these two alone. To illustrate the need for a third restriction, consider the following example and the following criterion.

Definition 1 (Absolute Number criterion (AN))**.**
For any two $F, F' \in \mathcal{F}$, we have $F' \succeq_{AN} F$ whenever

$$F_1'(y) \geq F_1(y) \qquad \textit{for all } y \in Y.$$

Moreover, $F' \succ_{AN} F$ if, in addition,

$$F_1'(y^*) > F_1(y^*) \qquad \textit{for some } y^* \in Y.$$

The Absolute Number criterion compares mechanisms based on their respective numbers of equilibria in X. This criterion satisfies our first two basic properties and its logic is implicitly used by Chen and Kesten (2017) (Theorem 2) when comparing the stability of school choice mechanisms.

To see why $\succeq_{AN}$ may be problematic, assume that there is a unique type profile y, which for two mechanisms $\tilde{F}$ and $\tilde{F}'$ is such that $\tilde{F}(y) = (1, 1)$ and $\tilde{F}'(y) = (4, 2)$. Clearly, the AN criterion concludes that $\tilde{F}'$ performs strictly better than $\tilde{F}$ because $\tilde{F}_1(y) = 1 < 2 = \tilde{F}_1'(y)$. This strict comparison is debatable because it ignores the fact that both mechanisms admit equilibria not in X and $\tilde{F}'$ admits more equilibria not in X than $\tilde{F}$ ($\tilde{F}_0(y) = 1 < 4 = \tilde{F}_0'(y)$). Even if $\tilde{F}'$ has twice as many equilibria in X as $\tilde{F}$, it not clear one should conclude that $\tilde{F}'$ performs *strictly* better than $\tilde{F}$ because $\tilde{F}'$ has four times as many equilibria not in X as $\tilde{F}$.

The issue with the AN criterion is that it violates a third basic property. *Replication Invariance* requires that two mechanisms that have the same proportion of their equilibria in X be viewed as performing equally well in terms of X. More precisely, if two mechanisms are equivalent on all but one type profile where one mechanism has k times as many equilibria in X and k times as many equilibria not in X as the other mechanism, then *Replication Invariance* concludes that the two mechanisms perform equally well in terms of X. When this is the case, we say that the former is a *k-replication* of the latter.

Axiom 4 (*Replication Invariance*)**.** *For all $F, F' \in \mathcal{F}$ and $k \in \mathbb{N}$, if (i) $F'_0(y^*) = kF_0(y^*)$ and $F'_1(y^*) = kF_1(y^*)$ for some y^*, and (ii) $F'(y) = F(y)$ for all $y \neq y^*$, then $F' \sim F$.*

It is easy to see how the axioms introduced thus far reach a different comparison of $\tilde{F} = (1, 1)$ and $\tilde{F}' = (4, 2)$ than $\succeq_{AN}$. Consider a third mechanism $\tilde{F}''$ such that $\tilde{F}''(y) = (2, 1)$. By *Monotonicity*, $\tilde{F}$ performs weakly better than $\tilde{F}''$. By *Replication Invariance*, because $\tilde{F}'$ has twice as many equilibria in X and twice as many equilibria not in X as $\tilde{F}''$, they perform equally well. Together, we must conclude that $\tilde{F}$ performs weakly better than $\tilde{F}'$, in contradiction with the comparison obtained with $\succeq_{AN}$. The debatable comparison obtained with $\succeq_{AN}$ follows from its violation of *Replication Invariance.*

1.4.1 The Proportion criterion

As we show in Theorem 1, these three axioms jointly characterize the *Proportion* criterion. It compares mechanisms based on the proportion of their equilibria in X.[6] This criterion does not come as a surprise given its reliance on *Replication Invariance*.

Definition 2 (Proportion criterion (PROP)).
For any two $F, F' \in \mathcal{F}$, we have $F' \succeq_{PROP} F$ whenever

$$\frac{F_1'(y)}{F_0'(y) + F_1'(y)} \geq \frac{F_1(y)}{F_0(y) + F_1(y)} \qquad \textit{for all } y \in Y.$$

Moreover, $F' \succ_{PROP} F$ if, in addition,

$$\frac{F_1'(y^*)}{F_0'(y^*) + F_1'(y^*)} = 1 \ \ \textit{and} \ \ \frac{F_1(y^*)}{F_0(y^*) + F_1(y^*)} = 0 \qquad \textit{for some } y^* \in Y.$$

Observe that, in line with *Domination*, the Proportion criterion yields strict comparisons only if there is a type profile where one mechanism has all its equilibria in X while all the equilibria of the other mechanism are not in X.

Our first result shows that the Proportion criterion is the *coarsest relation* satisfying our axioms.

Definition 3 (Coarsest relation).
A partial order $\succeq_{co}$ is the coarsest relation satisfying a set of axioms if

1. *$\succeq_{co}$ satisfies the set of axioms.*

[6]Under our assumptions, the proportion is always well-defined. Indeed, we assume that solution concepts admit at least one equilibrium for each type profile. As a result, the denominateur of the proportion is never zero.

2. For all $F, F' \in \mathcal{F}$ and all $\succeq$ satisfying the set of axioms,

$$F' \succeq_{co} F \Rightarrow F' \succeq F, \text{ and} \tag{1.1}$$

$$F' \succ_{co} F \Rightarrow F' \succ F. \tag{1.2}$$

A partial order that is the coarsest relation satisfying a set of axioms is not necessarily the only partial order that satisfies this set of axioms. Yet, the coarsest relation is the only partial order that satisfies the set of axioms while remaining silent on all pairs (of functions) that are not ranked by the joint implications of the axioms.

Theorem 1 identifies the close connection between the Proportion criterion and our three basic axioms.[7]

Theorem 1. *The partial order $\succeq_{PROP}$ is the coarsest relation satisfying Domination, Monotonicity and Replication Invariance.*

Proof.

Part 1 of Definition 3: The proof that $\succeq_{PROP}$ satisfies these three axioms is straightforward, and is therefore omitted.

[7]These three axioms are independent. Showing independence of *Monotonicity* is the most difficult part. We propose the criterion $I2$, which satisfies all these axioms except *Monotonicity*. Criterion $I2$ is based on the following function $f : [0,1] \to [0,1]$ defined as $f(x) = 1 - x$ for $x \in \{0,1\}$ and $f(x) = x$ for all $x \in (0,1)$. That is, function f is strictly increasing for all $x \in (0,1)$, but returns the smallest value for $x = 1$ and the greatest for $x = 0$. For any two $F, F' \in \mathcal{F}$, we have $F' \succeq_{I2} F$ whenever $f\left(\frac{F_1'(y)}{F_0'(y)+F_1'(y)}\right) \geq f\left(\frac{F_1(y)}{F_0(y)+F_1(y)}\right)$ for all $y \in Y$, and we have $F' \succ_{I2} F$ if in addition the inequality is strict for some $y^* \in Y$.

Implication (1.2) **in part 2 of Definition 3**: $F' \succ_{PROP} F \Rightarrow F' \succ F$

We slightly abuse notation and often write F and F instead of $F(y)$ and $F'(y)$ whenever there is no ambiguity on y. Let $Y^1 = \{y \in Y \mid F'_0 = 0 \text{ and } F_1 = 0\}$ be the set of type profiles for which all equilibria of F' are in X while all the equilibria of F are not in X. Since $F' \succ_{PROP} F$, we have that Y^1 is not empty and also that $\frac{F'_1}{F'_0+F'_1} \geq \frac{F_1}{F_0+F_1}$ for all $y \in Y$.

We show that any partial order $\succeq$ satisfying the list of axioms is such that $F' \succ F$ by constructing two sequences of functions $(L^p)_{p\in\{0,1\}}$ and $(K^p)_{p\in\{0,1\}}$ with $L^p, K^p \in \mathcal{F}$ such that

- $L^0 \succ K^0$,
- $L^1 \succeq L^0$ and $K^0 \succeq K^1$,
- $L^1 = F'$ and $K^1 = F$.

If these two sequences exist, then we have indeed that $F' \succ F$.

We construct each function in the sequence type profile by type profile. First, we construct L^0 and K^0. For all $y \in Y^1$, we take $L^0 = F'$ and $K^0 = F$. For all $y \in Y \backslash Y^1$, we take $L^0_1 = K^0_1 = (F'_0 + F'_1) * F_1$ and $L^0_0 = K^0_0 = (F_0 + F_1) * F'_0$. By successive applications of *Domination* we have $L^0 \succ K^0$. By "successive applications" of *Domination*, we mean that it is straightforward to construct a sequence of functions $(F^p)_{p\in\{0,\dots,P\}}$ with $F^0 = K^0$, $F^P = L^0$ and such that $F^{p+1} \succ F^p$ by the virtue of *Domination* for all $p \in \{0, \dots, P-1\}$.

Then, we construct L^1 and K^1 from L^0 and K^0 by changing their images on $Y\backslash Y^1$. For all $y \in Y^1$, we take $L^1 = L^0$ and

$K^1 = K^0$. For their construction on $Y \backslash Y^1$, we define two sequences $(\hat{L}^q)_{q\in\{0,1\}}$ and $(\hat{K}^q)_{q\in\{0,1\}}$ with $\hat{L}^q, \hat{K}^q \in \mathcal{F}$ such that

- $K^0 \succeq \hat{K}^0$,
- $\hat{L}^0 \succeq L^0$,
- $\hat{K}^1 \sim \hat{K}^0$,
- $\hat{L}^1 \sim \hat{L}^0$,

and we take $L^1 = \hat{L}^1$ and $K^1 = \hat{K}^1$, which implies $L^1 \succ K^1$. For any $y \in Y \backslash Y^1$, we take $\hat{L}^0_0 = L^0_0$ and $\hat{L}^0_1 = L^0_1 + (F'_1 * F_0 - F'_0 * F_1)$, where we have $F'_1 * F_0 - F'_0 * F_1 \geq 0$ because $\frac{F'_1}{F'_1+F'_0} \geq \frac{F_1}{F_1+F_0}$. We have $\hat{L}^0 \succeq L^0$ by (successive applications of) *Monotonicity*. For any $y \in Y \backslash Y^1$, we also take $\hat{K}^0_0 = K^0_0 + (F'_1 * F_0 - F'_0 * F_1)$ and $\hat{K}^0_1 = K^0_1$. We have $K^0 \succeq \hat{K}^0$ by (successive applications of) *Monotonicity*.
Then, we construct $\hat{L}^1$ from $\hat{L}^0$ and $\hat{K}^1$ from $\hat{K}^0$. For any $y \in Y \backslash Y^1$, let $\hat{L}^0$ be a $(F_0 + F_1)$-replication of $\hat{L}^1$ and $\hat{K}^0$ a $(F'_0 + F'_1)$-replication of $\hat{K}^1$ so that we have $\hat{L}^1 \sim \hat{K}^0$ and $\hat{K}^1 \sim \hat{L}^0$ by (successive applications of) *Replication Invariance*.

By construction, we have $L^1 = F'$ and $K^1 = F$ which completes the proof.

Implication (1.1) **in part 2 of Definition 3**: $F' \succeq_{PROP} F \Rightarrow F' \succeq F$

The proof can straightforwardly be adapted from the argument provided above, and is therefore omitted. ■

Theorem 1 calls for three remarks.

First, observe that one can also find complete orders satisfying this set of axioms.[8]

Second, Theorem 1 would still hold if we restrict ourselves to solution concepts with only one outcome per type profile, such as for instance the "truthfulness" solution concept. For this special case, the criteria must rank functions whose domain of images is $Z' = \{(z_0, z_1) \in \mathbb{N}_0^2 | z_0 + z_1 = 1\}$. All issues associated with having multiple equilibria are ruled out. For this special case, only *Domination* has bite because the remaining three axioms are trivially satisfied. Observe that the Proportion criterion would still yield a partial ranking of mechanisms. This illustrates that the difficulty to characterize a complete order is also present even when the equilibrium is unique.

Third, using the strict versions of axioms *Monotonicity*, i.e., if one mechanism has one more equilibrium in (resp. not in) X than the other mechanism, it performs strictly better (resp. worse) in terms of X, would lead to an impossibility because this stronger axiom is directly incompatible with *Replication Invariance*.

Although the Proportion criterion is very natural, it is af-

[8] Consider the following complete order. *For any two $F, F' \in \mathcal{F}$, we have $F' \succeq_{COMP} F$ whenever*

$$\sum_{y \in Y} \frac{F_1'(y)}{F_0'(y) + F_1'(y)} \geq \sum_{y \in Y} \frac{F_1(y)}{F_0(y) + F_1(y)}$$

Observe that these three axioms do not jointly imply this order. Additional properties would be required, typically imposing some form(s) of anonymity.

fected by two important limitations. Since the Proportion criterion relies on relatively weak axioms, it provides a very partial ranking and is thus often silent. Moreover, the Proportion criterion is not very workable. Indeed, this criterion requires computing the exact number of equilibria in each type profile, which can get quite challenging as even very simplified type profiles can admit multiple strategy profiles.

1.4.2 Dual extension of the Proportion criterion

To obtain more complete partial orders, we maintain the axioms imposed thus far while imposing additional restrictions that increase the number of pairs a partial order can compare. In this sense, our new partial orders extend the comparisons from $\succeq_{PROP}$ (they compare in the same way all pairs for which $\succeq_{PROP}$ makes an affirmative comparison, and reach affirmative comparisons for some pairs for which $\succeq_{PROP}$ is silent).

First, we consider an additional robustness axiom that we call *Consistency to Additional* $\in X$. Loosely put, *Consistency to Additional* $\in X$ requires that a comparison would not be altered if both mechanisms had one additional equilibrium in X. More precisely, assume that one mechanism M performs better than another M' (in terms of X). Consider slight variants of these two mechanisms such that, on a single type profile, both variants have one additional equilibrium in X. *Consistency to Additional* $\in X$ requires that the variant of M also performs better than the variant of M'.

Axiom 5 (*Consistency to Additional* $\in X$)**.** *For all* $F, F', \hat{F}, \hat{F}' \in \mathcal{F}$*, if (i)* $\hat{F}_0(y^*) = F_0(y^*)$*,* $\hat{F}'_0(y^*) = F'_0(y^*)$*,* $\hat{F}_1(y^*) = F_1(y^*)+1$

and $\hat{F}_1'(y^*) = F_1'(y^*) + 1$ *for some* y^**, and (ii)* $\hat{F}(y) = F(y)$ *and* $\hat{F}'(y) = F'(y)$ *for all* $y \neq y^*$*, then* $F' \succeq F \Rightarrow \hat{F}' \succeq \hat{F}$ *and* $F' \succ F \Rightarrow \hat{F}' \succ \hat{F}$*.*

Even if one may consider that *Consistency to Additional* $\in X$ is somewhat less compelling than our three basic axioms, we believe that it constitutes a plausible way of extending their affirmative comparisons. Observe in particular that *Consistency to Additional* $\in X$ does not impose any affirmative comparison on its own. It is only in combination with other axioms that it extends their pre-existing affirmative comparisons to more pairs.

Theorem 2 presented below shows that *Consistency to Additional* $\in X$ is exactly the difference between $\succeq_{PROP}$ and our second criterion. This criterion compares mechanisms by focusing exclusively on those type profiles for which all equilibria are in X or those for which all equilibria are not in X. More precisely, the criterion considers that a mechanism performs at least as well as another if the latter has no equilibria in X whenever the former has no equilibria in X *and* if the former has all its equilibria in X whenever the latter has all its equilibria in X. The comparison becomes strict if for some type profile, the former has all its equilibria in X whereas the latter has not.

Definition 4 (Profiles with Homogeneous Outcomes criterion (PHO))**.**
For any two $F, F' \in \mathcal{F}$*, we have* $F' \succeq_{PHO} F$ *if for all* $y \in Y$

$$F_1'(y) = 0 \Rightarrow F_1(y) = 0, \textit{and}$$
$$F_0(y) = 0 \Rightarrow F_0'(y) = 0.$$

Moreover, we have $F' \succ_{PHO} F$ if in addition

$$F_0'(y^*) = 0 \quad \textit{and} \quad F_0(y^*) > 0 \qquad \textit{for some } y^* \in Y.$$

The partial order $\succeq_{PHO}$ is more workable than $\succeq_{PROP}$. Indeed, obtaining affirmative comparisons with $\succeq_{PHO}$ never requires computing the proportion of equilibria in X. Even better, it is not even necessary to compute the number of equilibria associated to each type profile. The reason is that the affirmative comparisons of $\succeq_{PHO}$ are only based on type profiles for which all equilibria are equivalent in terms of X.

Observe that $\succeq_{PHO}$ yields strict comparisons only if there is a type profile where one mechanism has all its equilibria in X while the other's are *not all* in X. In contrast, in the case of $\succeq_{PROP}$, strict comparisons require that there is a type profile where one mechanism has all its equilibria in X while the other's are *all not* in X. Clearly, the weaker condition for strict comparisons under $\succeq_{PHO}$ derives from *Consistency to Additional* $\in X$, which extends the strict comparisons obtained from *Domination*.

Theorem 2 identifies the close connection between $\succeq_{PHO}$ and our four axioms.

Theorem 2. *The partial order $\succeq_{PHO}$ is the coarsest relation satisfying Domination, Monotonicity, Replication Invariance and Consistency to Additional $\in X$.*

Proof. **Part 1 of Definition 3**:

The PHO criterion clearly satisfies *Domination*, *Monotonicity* and *Replication Invariance.* We only prove that the PHO criterion satisfies *Consistency to Additional* $\in X$. We must show that, when its preconditions are met, we have $F' \succeq_{PHO} F \Rightarrow \hat{F}' \succeq_{PHO} \hat{F}$ and $F' \succ_{PHO} F \Rightarrow \hat{F}' \succ_{PHO} \hat{F}$. As the proof of the two implications are very similar, we only prove the latter.

Again, we slightly abuse notation and write F_0 and F_1 instead of $F_0(y)$ and $F_1(y)$. Given that $F' \succ_{PHO} F$, we can partition $Y = Y^1 \cup Y^2 \cup Y^3 \cup Y^4 \cup Y^5$, where

$$
\begin{aligned}
Y^1 &= \{y \in Y \mid F_0' = 0 \text{ and } F_1 = 0\},\\
Y^2 &= \{y \in Y \mid F_0' = 0 \text{ and } F_0 = 0\},\\
Y^3 &= \{y \in Y \mid F_0' = 0 \text{ and } F_0 > 0 \text{ and } F_1 > 0\},\\
Y^4 &= \{y \in Y \mid F_0' > 0 \text{ and } F_1' > 0 \text{ and } F_0 > 0 \text{ and } F_1 > 0\},\\
Y^5 &= \{y \in Y \mid F_0' > 0 \text{ and } F_1 = 0\},
\end{aligned}
$$

and where $Y^1 \cup Y^3$ is not empty. Such partition is illustrated in Figure 1.1. We show that when comparing $\hat{F}$ and $\hat{F}'$, we

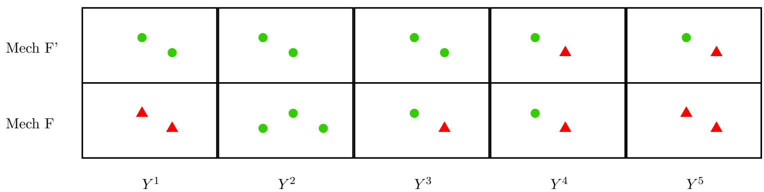

Figure 1.1: Illustration of type profiles for each subset of the partition.
Each green dot represents an equilibrium in X and each red triangle represents an equilibrium not in X.

can also partition $Y = \hat{Y}^1 \cup \hat{Y}^2 \cup \hat{Y}^3 \cup \hat{Y}^4 \cup \hat{Y}^5$ with the same definitions as above, except that these definitions consider functions $\hat{F}$ and $\hat{F}'$ instead of F and F', i.e., $\hat{Y}^1 = \{y \in Y \mid \hat{F}'_0 = 0 \text{ and } \hat{F}_1 = 0\}$, $\hat{Y}^2 = \{y \in Y \mid \hat{F}'_0 = 0 \text{ and } \hat{F}_0 = 0\}$, and so on. Moreover $\hat{Y}^1 \cup \hat{Y}^3$ is not empty. If we can partition Y in this way, then we have $\hat{F}' \succ_{PHO} \hat{F}$.

There remains to show that the preconditions of *Consistency to Additional* $\in X$, which link F and F' to $\hat{F}$ and $\hat{F}'$, are such that any $y \in Y^1 \cup Y^2 \cup Y^3 \cup Y^4 \cup Y^5$ is such that $y \in \hat{Y}^1 \cup \hat{Y}^2 \cup \hat{Y}^3 \cup \hat{Y}^4 \cup \hat{Y}^5$ and any $y \in Y^1 \cup Y^3$ is such that $y \in \hat{Y}^1 \cup \hat{Y}^3$. For all $y \neq y^*$, we have $\hat{F}(y) = F(y)$ and $\hat{F}'(y) = F'(y)$, which directly implies that for all $p \in \{1, \ldots, 5\}$ we have $y \in \hat{Y}^p$ when $y \in Y^p$. For y^*, we have $\hat{F}_0(y^*) = F_0(y^*)$ and $\hat{F}'_0(y^*) = F'_0(y^*)$, as well as $\hat{F}_1(y^*) = F_1(y^*) + 1$ and $\hat{F}'_1(y^*) = F'_1(y^*) + 1$. These preconditions are such that $y^* \in Y^1 \Rightarrow y^* \in \hat{Y}^3$, $y^* \in Y^2 \Rightarrow y^* \in \hat{Y}^2$, $y^* \in Y^3 \Rightarrow y^* \in \hat{Y}^3$, $y^* \in Y^4 \Rightarrow y^* \in \hat{Y}^4$ and $y^* \in Y^5 \Rightarrow y^* \in \hat{Y}^4$. Finally, as $Y^1 \cup Y^3$ is non-empty, $y^* \in Y^1 \Rightarrow y^* \in \hat{Y}^3$ and $y^* \in Y^3 \Rightarrow y^* \in \hat{Y}^3$, $\hat{Y}^1 \cup \hat{Y}^3$ is not empty, the desired result.

Implication (1.2) **in part 2 of Definition 3**: $F' \succ_{PHO} F \Rightarrow F' \succ F$

Since $F' \succ_{PHO} F$, we can partition $Y = Y^1 \cup Y^2 \cup Y^3 \cup Y^4 \cup Y^5$ using the same definitions used for part 1, and moreover $Y^1 \cup Y^3$ is not empty.

We show that any partial order $\succeq$ satisfying the list of axioms is such that $F' \succ F$ by constructing two sequences of functions $(L^p)_{p \in \{0,\ldots,5\}}$ and $(K^p)_{p \in \{0,\ldots,5\}}$ with $L^p, K^p \in \mathcal{F}$ such that

- $L^0 \succ K^0$,
- $L^{p+1} \succeq L^p$ and $K^p \succeq K^{p+1}$ for all $p \in \{0, \ldots, 4\}$,

- $L^5 = F'$ and $K^5 = F$.

If such two sequences exist, then we have indeed that $F' \succ F$.

First, we define L^0 and K^0. We construct these two functions type profile by type profile. For all $y \in Y^1 \cup Y^3$ we take $L_0^0 = K_1^0 = 0$ and $L_1^0 = K_0^0 = 1$. For all $y \in Y^2$ we take $L_0^0 = K_0^0 = 0$ and $L_1^0 = K_1^0 = 1$. For all $y \in Y^4 \cup Y^5$ we take $L_0^0 = K_0^0 = 1$ and $L_1^0 = K_1^0 = 0$. By (successive applications of) *Domination*, we have $L^0 \succ K^0$.

We define the remaining elements of the two sequences in 5 successive steps, one for each subset in the partition of Y. Functions L^p and K^p are constructed from L^{p-1} and K^{p-1} in step p in such a way that for all $a \in \{1, \ldots, p\}$ and all $y \in Y^a$ we have $L^p(y) = F'(y)$ and $K^p(y) = F(y)$. When the construction of a function is left unspecified on a type profile, it means that this function takes the same image as the function from which it is constructed.

- **Step 1**: Define L^1 and K^1 from L^0 and K^0 by changing their images on Y^1.

  For any $y \in Y^1$, take $L_0^1 = K_1^1 = 0$ and $L_1^1 = F_1'$ and $K_0^1 = F_0$. That is, $L^1(y)$ is a F_1'-replication of $L^0(y)$ and $K^1(y)$ is a F_0-replication of $K^0(y)$. By (successive applications of) *Replication Invariance*, we have $L^1 \sim L^0$ and $K^1 \sim K^0$.

- **Step 2**: Define L^2 and K^2 from L^1 and K^1 by changing their images on Y^2.

  For any $y \in Y^2$, take $L_0^2 = K_0^2 = 0$ and $L_1^2 = F_1'$ and $K_0^2 = F_1$. That is, $L^2(y)$ is a F_1'-replication of $L^1(y)$ and $K^2(y)$ is a F_1-replication of $K^1(y)$. By (successive applications of) *Replication Invariance*, we have $L^2 \sim L^1$ and $K^2 \sim K^1$.

- **Step 3**: Define L^3 and K^3 from L^2 and K^2 by changing their images on Y^3.

  We define two sequences $(\hat{L}^q)_{q\in\{0,1\}}$ and $(\hat{K}^q)_{q\in\{0,1\}}$ with

  - $\hat{K}^0 \sim K^2$,
  - $\hat{L}^0 \succ \hat{K}^1$,
  - $\hat{L}^1 \sim \hat{L}^0$,

  and we take $L^3 = \hat{L}^1$ and $K^3 = \hat{K}^1$, which implies $L^3 \succ K^3$.

  For any $y \in Y^3$, we take $\hat{K}^0_0 = F_0$ and $\hat{K}^0_1 = 0$. As $\hat{K}^0(y)$ is a F_0-replication of $K^2(y)$, by (successive applications of) *Replication Invariance*, we have $\hat{K}^0 \sim K^2$. We then construct $\hat{L}^0$ from L^2 and $\hat{K}^1$ from $\hat{K}^0$ by addition of the same number of equilibria in $X(y)$. For any $y \in Y^3$, we take $\hat{L}^0_0 = L^2_0$, $\hat{K}^1_0 = \hat{K}^0_0$, $\hat{L}^0_1 = L^2_1 + F_1$ and $\hat{K}^1_1 = \hat{K}^0_1 + F_1$. By transitivity we have from the previous steps that $L^2 \succ \hat{K}^0$. Therefore, we get $\hat{L}^0 \succ \hat{K}^1$ by (successive applications of) *Consistency to Additional* $\in X$. For any $y \in Y^3$, we take $\hat{L}^1_0 = 0$ and $\hat{L}^1_1 = F'_1$. As $\hat{L}^1(y)$ is a $\frac{F'_1}{1+F_1}$-replication of $\hat{L}^0(y)$, by (successive applications of) *Replication Invariance*, we have $\hat{L}^1 \sim \hat{L}^0$. If $\frac{F'_1}{1+F_1}$ is not an integer, then an intermediary function $\hat{L}^*$ must be defined such that $\hat{L}^*(y)$ is a F'_1-replication of $\hat{L}^0(y)$ and, also, such that $\hat{L}^*(y)$ is a $(1 + F_1)$-replication of $\hat{L}^1(y)$.

- **Step 4**: Define L^4 and K^4 from L^3 and K^3 by changing their images on Y^4.

We define two sequences $(\hat{L}^q)_{q\in\{0,\dots,3\}}$ and $(\hat{K}^q)_{q\in\{0,1\}}$ with

- $\hat{L}^0 \sim L^3$ and $\hat{K}^0 \sim K^3$,
- $\hat{L}^1 \succ \hat{K}^1$,
- $\hat{L}^2 \succeq \hat{L}^1$ and $\hat{L}^3 \sim \hat{L}^2$,

and we take $L^4 = \hat{L}^3$ and $K^4 = \hat{K}^1$, which implies $L^4 \succ K^4$.

For any $y \in Y^4$, we take $\hat{L}^0_0 = F'_0 F_1$, $\hat{K}^0_0 = F_0$ and $\hat{L}^0_1 = \hat{K}^0_1 = 0$. As $\hat{L}^0(y)$ is a $F'_0 F_1$-replication of $L^3(y)$ and $\hat{K}^0(y)$ is a F_0-replication of $K^3(y)$, by (successive applications of) *Replication Invariance*, we have $\hat{L}^0 \sim L^3$ and $\hat{K}^0 \sim K^3$. We then construct $\hat{L}^1$ from $\hat{L}^0$ and $\hat{K}^1$ from $\hat{K}^0$ by addition of the same number of equilibria in $X(y)$. For any $y \in Y^4$, we take $\hat{L}^1_0 = \hat{L}^0_0$, $\hat{K}^1_0 = \hat{K}^0_0$, $\hat{L}^1_1 = \hat{L}^0_1 + F_1$ and $\hat{K}^1_1 = \hat{K}^0_1 + F_1$. By transitivity we have from the previous steps that $\hat{L}^0 \succ \hat{K}^0$. Therefore, we get $\hat{L}^1 \succ \hat{K}^1$ by (successive applications of) *Consistency to Additional* $\in X$. For any $y \in Y^4$, we take $\hat{L}^2_0 = \hat{L}^1_0$ and $\hat{L}^2_1 = \hat{L}^1_1 + F_1(F'_1 - 1)$. By (successive applications of) *Monotonicity*, we have $\hat{L}^2 \succeq \hat{L}^1$. Finally, for any $y \in Y^4$, we take $\hat{L}^3_0 = F'_0$ and $\hat{L}^3_1 = F'_1$. As $\hat{L}^2(y)$ is a F_1-replication of $\hat{L}^3(y)$, by (successive applications of) *Replication Invariance*, we have $\hat{L}^3 \sim \hat{L}^2$.

- **Step 5**: Define L^5 and K^5 from L^4 and K^4 by changing their images on Y^5.

  We define a sequence $(\hat{L}^q)_{q\in\{0,1\}}$ and a function $\hat{K}^0$ with

  - $\hat{L}^0 \sim L^4$ and $\hat{K}^0 \sim K^4$,

– $\hat{L}^1 \succeq \hat{L}^0$,

and we take $L^5 = \hat{L}^1$ and $K^5 = \hat{K}^0$, which implies $L^5 \succ K^5$.

For any $y \in Y^5$, we take $\hat{L}^0_0 = F'_0$, $\hat{K}^0_0 = F_0$ and $\hat{L}^0_1 = \hat{K}^0_1 = 0$. As $\hat{L}^0(y)$ is a F'_0-replication of $L^4(y)$ and $\hat{K}^0(y)$ is a F_0-replication of $K^4(y)$, by (successive applications of) *Replication Invariance*, we have $\hat{L}^0 \sim L^4$ and $\hat{K}^0 \sim K^4$. For any $y \in Y^5$, we take $\hat{L}^1_0 = \hat{L}^0_0$ and $\hat{L}^1_1 = \hat{L}^0_1 + F'_1$. By (successive applications of) *Monotonicity*, we have $\hat{L}^1 \succeq \hat{L}^0$.

By construction, we have $L^5 = F'$ and $K^5 = F$, which completes the proof.

Implication (1.1) **in part 2 of Definition 3**:

The proof can straightforwardly be adapted from the argument provided above, and is therefore omitted.

■

Interestingly, imposing *Consistency to Additional* $\in X$ allows comparing two mechanisms by only focusing on the subset of type profiles for which all equilibria are equivalent in terms of X. In a sense, type profiles for which both mechanisms yield some outcomes in X and some outcomes not in X are "irrelevant" for the comparison.[9] This greatly increases the number of pairs that can be compared because, unlike $\succeq_{PROP}$, the

[9]These type profiles are irrelevant in the sense that the exact fraction of outcomes in X of each mechanism does not matter.

partial order $\succeq_{PHO}$ is not necessarily bogged down by the existence of a few type profiles for which proportions of equilibria in X are reversed.

This reduction of the domain of "relevant" type profile is rather surprising. We emphasize that *Consistency to Additional* $\in X$ is not sufficient by itself to yield such reduction. In fact, the Absolute Number criterion satisfies *Consistency to Additional* $\in X$ together with *Domination* and *Monotonicity*, but still bases its comparisons on all type profiles in the domain. This reduction is the result of the combination of the list of axioms used.

Finally, we show that an axiom that is dual to *Consistency to Additional* $\in X$ leads to a criterion that is dual to $\succeq_{PHO}$. *Consistency to Additional* $\notin X$ preserves the logic of *Consistency to Additional* $\in X$, but the former focuses on equilibria not in X, whereas the latter focuses on equilibria in X. More precisely, assume that one mechanism M performs better than another M' (in terms of X). Consider slight variants of these two mechanisms such that, on a single type profile, both variants have one additional equilibrium *not* in X. *Consistency to Additional* $\notin X$ requires that the variant of M also performs better than the variant of M'.

Axiom 6 (*Consistency to Additional* $\notin X$)**.** *For all* $F, F', \hat{F}, \hat{F}' \in \mathcal{F}$*, if (i)* $\hat{F}_0(y^*) = F_0(y^*) + 1$*,* $\hat{F}'_0(y^*) = F'_0(y^*) + 1$*,* $\hat{F}_1(y^*) = F_1(y^*)$ *and* $\hat{F}'_1(y^*) = F'_1(y^*)$ *for some* y^**, and (ii)* $\hat{F}(y) = F(y)$ *and* $\hat{F}'(y) = F'(y)$ *for all* $y \neq y^*$*, then* $F' \succeq F \Rightarrow \hat{F}' \succeq \hat{F}$ *and* $F' \succ F \Rightarrow \hat{F}' \succ \hat{F}$*.*

Unsurprisingly, the partial order $\succeq_{PHO^*}$ associated to *Consistency to Additional* $\notin X$ is very similar to $\succeq_{PHO}$. In fact, the weak comparisons of these two criteria are based on the same conditions. The difference comes from the condition for strict comparisons. The partial order $\succeq_{PHO^*}$ yields strict comparisons only if there is a type profile where one mechanism has *some* of its equilibria in X while the other has *none* of its equilibria in X.

Definition 5 (Profiles with Homogeneous Outcomes criterion* (PHO^*))**.**
For any two $F, F' \in \mathcal{F}$, we have $F' \succeq_{PHO^*} F$ if for all $y \in Y$

$$\begin{aligned} F_1'(y) = 0 &\Rightarrow F_1(y) = 0, \textit{and} \\ F_0(y) = 0 &\Rightarrow F_0'(y) = 0. \end{aligned}$$

Moreover, we have $F' \succ_{PHO^*} F$ if in addition

$$F_1'(y^*) > 0 \quad \textit{and} \quad F_1(y^*) = 0 \qquad \textit{for some } y^* \in Y.$$

Theorem 3 identifies the close connection between $\succeq_{PHO^*}$ and the four axioms.

Theorem 3. *The partial order $\succeq_{PHO^*}$ is the coarsest relation satisfying Domination, Monotonicity, Replication Invariance and Consistency to Additional $\notin X$.*

Proof. The proof can straightforwardly be adapted from the proof of Theorem 2, and is therefore omitted. ■

1.5 Illustration with school choice mechanisms

For illustrative purposes, we compare two matching mechanisms for the allocation of school seats. In this context, a matching algorithm determines the allocation of seats based on the preferences reported by the students and the priorities that students receive at the different schools. The players are the students and their strategy set is the set of preferences they can report. A type profile consists in a preference profile together with a priority profile. The complete description of the school choice model considered is given in Appendix 1.A.

We focus on an extremely simplified domain of school choice problems, with only three students and three schools, each endowed with one seat. On this narrow domain, we compare two school choice mechanisms with respect to the “stable" social choice correspondence, which is central in the school choice literature. This fairness property essentially requires that no blocking pair exists in the assignment.[10] The two mechanisms we compare are constrained versions of the Deferred Acceptance (DA) and Boston (BOS) mechanisms, for which students are allowed to report preferences on *two* schools only (Haeringer and Klijn (2009)). We denote these mechanisms as DA^2 and BOS^2.[11] For both mechanisms, we use *undominated strategy profile* as a solution concept. Although it is widely used, the Nash equilibrium solution concept might not

[10]An assignment has a blocking pair if a student is assigned to a school that another student prefers to her assignment and the other student has higher priority at this school than the first student.

[11]See Appendix 1.B for the description of both mechanisms.

be credible for such mechanisms. In school choice, Nash equilibrium may require a degree of coordination that goes beyond what can reasonably be expected from parents who play the corresponding game, often as a one-shot game. Experimental evidence also suggest that Nash equilibria are rarely reached in these mechanisms (Calsamiglia et al. (2010)).

Under both DA^2 and BOS^2, many type profiles admit multiple undominated strategy profiles, some of which lead to stable assignments while others do not. However, there are reasons to believe that DA^2 should be deemed more stable than BOS^2. First, theoretical results have shown that *un*constrained DA is stable in dominant strategies, whereas *un*constrained BOS is stable only in Nash equilibrium, i.e., when assuming complete coordination among the players. Second, experimental evidence shows that constrained versions of DA are more stable than constrained versions of BOS. In a constrained environment, i.e., when players can report preferences on a limited number of schools, Calsamiglia et al. (2010) show that, even though stable assignments rarely occur, there are significantly more blocking pairs arising in constrained versions of BOS than in constrained versions of DA. (Recall that the "stable" correspondence essentially selects assignments that do not contain any blocking pairs.) Also, Klijn et al. (2013) show that, independently of players' risk aversion, BOS is less likely to produce stable assignments than DA.

Unfortunately, Proposition 1 shows that the Proportion criterion cannot compare the stability of these two mechanisms: There exist type profiles for which the proportion of stable equilibria is greater under DA^2 than under BOS^2, as well as other type profiles for which the converse is true.

Proposition 1. *Let the solution concept C be undominated strategy profiles. Let X denote the stable correspondence. Let F^{DA^2} and F^{BOS^2} be the functions respectively associated to DA^2 and BOS^2 by C and X. There exists a type profile $y \in Y$ such that*

$$\frac{F_1^{DA^2}(y)}{F_0^{DA^2}(y) + F_1^{DA^2}(y)} = 1 \text{ and } \frac{F_1^{BOS^2}(y)}{F_0^{BOS^2}(y) + F_1^{BOS^2}(y)} < 1,$$

and a type profile $y' \in Y$ such that

$$\frac{F_1^{DA^2}(y')}{F_0^{DA^2}(y') + F_1^{DA^2}(y')} < \frac{F_1^{BOS^2}(y')}{F_0^{BOS^2}(y') + F_1^{BOS^2}(y')}.$$

Proof. Type profile y is presented in Table 1.1. For visual convenience, the schools at which a student has top-priority are starred.

R_{i_1}	R_{i_2}	R_{i_3}
s_1^*	s_1	s_2
s_2	s_2^*	s_3^*
s_3	s_3	s_1

$\trianglerighteq_{s_1}$	$\trianglerighteq_{s_2}$	$\trianglerighteq_{s_3}$
i_1	i_2	i_3
$\vdots$	$\vdots$	$\vdots$

Table 1.1: Type profile y for which all undominated strategy profiles under DA^2 are stable but not all undominated strategy profiles under BOS^2.

First, we show for y that all undominated strategy profiles under DA^2 are stable. As i_1 has a top-priority at her most-preferred school, she has a dominant strategy and is assigned to that school under any undominated strategy profile (Lemma 3). As i_2 and i_3 have a top-priority at their second most-preferred school, they have a dominant strategy (Lemma 4) that ranks their two most-preferred schools according to their true preference (Lemma 5). Therefore, only one assignment can be reached under undominated strategy profiles. This assignment is such that each student is assigned to her most-preferred top-priority school. This assignment is stable.

Second, we show for y that one undominated strategy profiles under BOS^2 is not stable. Consider the reported profile Q shown herebelow.

$$\begin{aligned} Q_{i_1} &: \; s_1 \; s_2 \\ Q_{i_2} &: \; s_1 \; s_2 \\ Q_{i_3} &: \; s_2 \; s_3 \end{aligned}$$

Profile Q is an undominated strategy profile under BOS^2 (Lemma 6). The assignment $BOS^2(Q)$ is such that i_1 is assigned to s_1, i_2 is unassigned and i_3 is assigned to s_2. This assignment is unstable as i_2 prefers s_2 over being unassigned and i_2 has a higher priority at s_2 than i_1.

Type profile y' is presented in Table 1.2. For visual convenience, the schools at which a student has top-priority are starred and the only stable assignment is boxed.

Under DA^2, one-third of undominated strategy outcomes are stable. It is a dominant strategy for both i_1 and i_2 to truthfully report their preference because they each have a top-priority at their second favorite school (Lemma 3). In turn, student i_3 has three undominated strategies (Lemma 5):

R_{i_1}	R_{i_2}	R_{i_3}
s_2	s_1	s_1
$\boxed{s_1^*}$	$\boxed{s_2^*}$	s_2
i_1	i_2	$\boxed{s_3}$

$\trianglerighteq_{s_1}$	$\trianglerighteq_{s_2}$	$\trianglerighteq_{s_3}$
i_1	i_2	i_1
i_3	$\vdots$	$\vdots$
i_2		

Table 1.2: Type profile y' for which the proportion of stable undominated strategy outcomes is larger under BOS^2 than under DA^2.

$$\begin{array}{rl} Q_{i_3}: & s_1 \ s_2 \\ Q'_{i_3}: & s_1 \ s_3 \\ Q''_{i_3}: & s_2 \ s_3 \end{array}$$

If i_3 reports Q_{i_3}, then the assignment is unstable because i_3 is unassigned while the seat at s_3 is vacant. If i_3 reports Q''_{i_3}, then the assignment is again unstable because i_2 is assigned to s_1 even if i_2 has a lower priority at s_1 than i_3. If i_3 reports Q'_{i_3}, then the assignment is stable.

Under BOS^2, more than one-third of undominated strategy outcomes are stable. Students i_1 and i_2 have two undominated strategies whereas i_3 has six undominated strategies (Lemma 6):

$$Q_{i_1}: \; s_2 \; s_1^* \qquad Q_{i_2}: \; s_1 \; s_2^* \qquad Q_{i_3}: \; s_1 \; s_2$$
$$Q'_{i_1}: \; s_1^* \; s_2 \qquad Q'_{i_2}: \; s_2^* \; s_1 \qquad Q'_{i_3}: \; s_1 \; s_3$$
$$Q''_{i_3}: \; s_2 \; s_3$$
$$Q'''_{i_3}: \; s_2 \; s_1$$
$$Q''''_{i_3}: \; s_3 \; s_1$$
$$Q'''''_{i_3}: \; s_3 \; s_2$$

We show that a proportion 10/24 of BOS^2 assignments are stable, which is larger than the proportion 1/3 obtained under DA^2.

First, we consider the six undominated strategy profiles for which i_1 and i_2 report Q_{i_1} and Q_{i_2}. None of the six assignments are stable, because for all of them we have either that i_3 is unassigned or i_2 is assigned to s_1.

Second, we consider the six undominated strategy profiles for which i_1 and i_2 report Q'_{i_1} and Q'_{i_2}. Under these profiles, i_1 is assigned to s_1 and i_2 is assigned to s_2. The assignment is stable if i_3 reports s_3, which is the case in all her undominated strategies but Q_{i_3} and Q'''_{i_3}. Hence, four out of these six assignments are stable.

Third, we consider the six undominated strategy profiles for which i_1 and i_2 report Q_{i_1} and Q'_{i_2}. Under these profiles, i_2 is assigned to s_2. The assignment is stable if i_3 reports s_3 and does not report s_1 first. Hence, three out of these six assignments are stable.

Fourth, we consider the six undominated strategy profiles for which i_1 and i_2 report Q'_{i_1} and Q_{i_2}. Under these profiles, i_1 is assigned to s_1. The assignment is stable if i_3 reports s_3 and does not report s_2 first. Hence, three out of these six assignments are stable.

■

The second important limitation of $\succeq_{PROP}$ is that this criterion is not very workable. Even in our extremely simplified domain with only three students and three schools, computing the exact number of equilibria in each type profile and identifying the proportion of these equilibria that are in X can be challenging. As we show in the proof of Proposition 1, the relatively simple type profile y' admits 24 different undominated strategy profiles under BOS^2. Investigating the stability of all 24 is quite cumbersome. What is more, the proof only shows that the two mechanisms cannot be compared by $\succeq_{PROP}$, which requires considering only two type profiles.

This example illustrates the need for partial orders that are less partial and more workable than $\succeq_{PROP}$.

The increased discriminatory power of $\succeq_{PHO}$ provides a sense for which we can affirmatively compare DA^2 and BOS^2 in terms of stability. This comparison is in line with our expectations.

Proposition 2. *Let the solution concept C be undominated strategy profiles. Let X denote the stable assignments correspondence. Letting F^{DA^2} and F^{BOS^2} be the functions respectively associated to DA^2 and BOS^2 by C and X, we have $F^{DA^2} \succ_{PHO} F^{BOS^2}$.*[12]

[12]Note that with F^{DA} and F^{BOS} the functions respectively associated to unrestricted DA and BOS by C and X, we also have $F^{DA} \succ_{PHO} F^{BOS}$. Indeed, in DA all students have a single dominant strategy which consist in ranking all their acceptable schools without switches. There is therefore only one undominated strategy profile in DA, and this profile is always stable. It is then sufficient to show that some of the many undominated

Proof. **Part 1.** $F^{DA^2} \succeq_{PHO} F^{BOS^2}$.

First, we show that for all $y \in Y$ for which no undominated strategy profiles under DA^2 leads to a stable assignment, no undominated strategy profiles under BOS^2 leads to a stable assignment. To do so, we show that there exists no $y \in Y$ for which no undominated strategy profiles under DA^2 leads to a stable assignment. Consider the contradiction assumption that, for some type profile $y^* \in Y$, no undominated strategy profile under DA^2 leads to a stable assignment. Let μ^* denote the most-efficient stable assignment for type profile y^*.

Assume first that all students are assigned to a school under μ^*. Consider any undominated strategy profile $Q = (Q_{i_1}, Q_{i_2}, Q_{i_3})$ under DA^2 for which each student i reports $\mu^*(i)$, the school to which she is assigned under μ^*. Any student i has an undominated strategy with this property. Indeed, if $\mu^*(i)$ is not her third favorite acceptable school, then reporting her two favorite acceptable schools in the order of her truthful preference is clearly undominated under DA^2. If $\mu^*(i)$ is her third favorite acceptable school, then i has no dominant strategy under DA^2 and reporting any two acceptable schools in the same order as the order of preference is undominated (Lemma 5).

Since strategies in Q are undominated, they report the schools in the order of the students' truthful preference (Lemma 5). Hence, if $\mu^*(i)$ is not reported first in Q_i, then i prefers the school reported first in Q_i over $\mu^*(i)$. Then, because μ^* is a stable assignment, either the assignment $DA^2(Q)$ is μ^*, which violates the contradiction assumption, or $DA^2(Q)$ is a Pareto improvement over μ^*. In the latter case, $DA^2(Q)$ is unstable

---

strategy profiles in BOS are not stable.

because μ^* is the most-efficient stable assignment. As $DA^2(Q)$ is an unstable Pareto improvement over μ^*, we have that one student, say i_3, is assigned under $DA^2(Q)$ to the same school as under μ^*, while i_1 and i_2 have exchanged the schools they are assigned to under μ^*. If all students are assigned to a different school as under μ^*, then μ^* cannot be the most-efficient stable assignment. This implies that $Q_{i_1} : \mu^*(i_2)\ \mu^*(i_1)$ and $Q_{i_2} : \mu^*(i_1)\ \mu^*(i_2)$. Assignment $DA^2(Q)$ is unstable because there is a school $s \in \{\mu^*(i_1), \mu^*(i_2)\}$ that i_3 prefers over $\mu^*(i_3)$ and i_3 has a higher priority at s than the student assigned to s under $DA^2(Q)$. As i_3 prefers s over $\mu^*(i_3)$ we have that $Q'_{i_3} :\ s\ \mu^*(i_3)$ is undominated under DA^2 (Lemma 5). As i_3 has a higher priority at s than the student assigned to s under $DA^2(Q)$, we must have that $DA^2(Q_{i_1}, Q_{i_2}, Q'_{i_3}) = \mu^*$, which violates the contradiction assumption.

Assume then that some student i is not assigned to a school under μ^*. Because there are three schools and three students, this implies that student i finds at most two schools acceptable. In turn, this implies that any student i' who is assigned to a school under μ^* is assigned either to her most-preferred school or to her second most-preferred school. The reason is that i is rejected from all of her acceptable schools. Hence, any school s that is acceptable for i is assigned under μ^* to another student i'. This is only possible if i' prefers s to the school that has a vacant seat under μ^*. Hence, any such student i' is assigned to a school she prefers to at least one other school.

Consider any strategy profile $Q = (Q_{i_1}, Q_{i_2}, Q_{i_3})$ under DA^2 for which each student reports either her only acceptable school, or her two most-preferred acceptable schools in the same order as the order of her true preference. All strategies in Q are undominated under DA^2 (Lemma 5). The contradiction assumption is violated because we have $DA^2(Q) = \mu^*$. Indeed, under μ^*, unassigned students find at most two schools acceptable and other students are assigned either to their most-preferred or their second most-preferred acceptable school. Therefore, on this type profile, the Deferred Acceptance mechanism stops before reaching the acceptable schools not reported in Q (if any). Hence, when the profile is Q, mechanism DA^2 follows the same steps as the Deferred Acceptance, and thus yields the most efficient stable assignment.

There remains to show that, for all $y \in Y$ for which all undominated strategy profiles under BOS^2 lead to a stable assignment, all undominated strategy profiles under DA^2 also lead to a stable assignment. The proof is based on Lemma 1, which shows that the set of assignments obtained by undominated strategy profiles under DA^2 are nested in the set of assignments obtained by undominated strategy profiles under BOS^2.

Lemma 1. *For any undominated strategy profile Q of DA^2, there exists an undominated strategy profile Q' of BOS^2 such that $DA^2(Q) = BOS^2(Q')$.*

Proof. Take any profile Q that is undominated under DA^2. Let assignment $\mu = DA^2(Q)$.

We construct a strategy profile Q' that is undominated under BOS^2 and such that $BOS^2(Q') = \mu$. For any student i who is *un*assigned under μ we let $Q'_i = Q_i$. For any student i who is assigned to a school under μ,

- we let $Q'_i = Q_i$ if $\mu(i)$ is reported first in Q_i,

- else Q'_i reports $\mu(i)$ first and also reports her most-preferred acceptable school different from $\mu(i)$ (if any).

First, we show that Q'_i is undominated under BOS^2. If i finds only one school acceptable, then reporting this school only is a dominant strategy under both BOS^2 and DA^2 (Lemma 2) and by construction this case is such that $Q'_i = Q_i$. Assume then that i finds at least two schools acceptable.

If the most-preferred school of student i is a top-priority school for i, then it is a dominant strategy to report this school first under DA^2 (Lemma 3) and i must be assigned to this school under μ, i.e., this school is $\mu(i)$. By construction, Q'_i reports $\mu(i)$ first, and therefore Q'_i is a dominant strategy under BOS^2 (Lemma 3). Assume then that i finds at least two schools acceptable and her most-preferred school is not a top-priority school for i.

- Case 1: $Q'_i = Q_i$.
  Since Q_i is undominated under DA^2, we have by Lemma 5 that Q_i reports two schools, ranks these two schools according to i's true preference and i weakly prefers these two schools over her most-preferred top-priority school. As $Q'_i = Q_i$, we then have that Q'_i is undominated under BOS^2 (Lemma 6).

- Case 2: $Q'_i \neq Q_i$.
  By construction of Q'_i, this case is such that i is assigned under μ and $\mu(i)$ is reported second in Q_i. Then, since Q_i is undominated under DA^2, by Lemma 5, i weakly prefers the two schools reported in Q_i over her most-preferred top-priority school. If $\mu(i)$ is i's most-preferred top-priority school, then Q'_i is undominated under BOS^2

(Lemma 6), because by construction of Q'_i this school is reported first in Q'_i. If $\mu(i)$ is not i's most-preferred top-priority school, then Q'_i is undominated under BOS^2 (Lemma 6) because by construction of Q'_i this strategy reports two schools, one of them being preferred to i's most-preferred top-priority school. (The school reported first in Q_i is strictly preferred to $\mu(i)$.)

Second, we show that $BOS^2(Q') = \mu$. Consider the subset I' of students who are unassigned under μ. Since $DA^2(Q) = \mu$, this implies that no student $i \in I'$ can be blocking in matching μ at a school she reports in Q_i. (Indeed, if such student i was blocking at a school s, then the student j for whom $\mu(j) = s$ should have been rejected from s in the course of DA^2 under Q, a contradiction.) In other words, the seat at the schools that i reports in Q_i are assigned under μ to competitors of i at these schools. By construction, for any student $i \in I'$ we have $Q'_i = Q_i$. Since any student $j \notin I'$ reports $\mu(j)$ first in Q'_j, this implies that all seats at all schools reported by any student $i \in I$ are assigned to competitors of i in the first round of BOS^2 under Q'. As a result, all students in I' are also unassigned under $BOS^2(Q')$. Finally, since any student $j \notin I'$ reports $\mu(j)$ first in Q'_j, student j is also assigned to $\mu(j)$ under $BOS^2(Q')$. Together, we have $BOS^2(Q') = \mu$.

■

Consider any $y \in Y$ for which all undominated strategy profiles under BOS^2 lead to a stable assignment. By Lemma 1, for any undominated strategy profiles under DA^2, there is an undominated strategy profiles under BOS^2 that leads to the same assignment. As a result, any undominated strategy profiles under DA^2 leads to a stable assignment.

Part 2. For some $y^* \in Y$, all undominated strategy profiles under DA^2 lead to stable assignments, while some undominated strategy profiles under BOS^2 leads to unstable assignments.

As shown in the proof of Proposition 1, type profile y presented in Table 1.1 has the required properties. Together, Part 1 and Part 2 imply that $F^{DA^2} \succ_{PHO} F^{BOS^2}$. ∎

The proof of Proposition 2 illustrates another reason why the partial order $\succeq_{PHO}$ is more workable than $\succeq_{PROP}$. Affirmative comparisons of $\succeq_{PHO}$ are only based on type profiles for which all equilibria are equivalent in terms of X. Importantly, it is sometimes easier to compare mechanisms on these particular type profiles. For instance, the proof of Proposition 2 takes advantage of the focus on these type profiles. A key step in the proof of Proposition 2 is that the US-assignments under DA^2 are nested in the US-assignments under BOS^2. (For short terminology, we refer to an assignment sustained by an undominated strategy profile under mechanism M simply as an US-assignment under M.) This directly implies that, if all US-assignments are stable under BOS^2, then all US-assignments are stable under DA^2. The weak comparison $F^{DA^2} \succeq_{PHO} F^{BOS^2}$ then follows from the fact that there is no type profile in our domain for which all US-assignments under DA^2 are unstable.

Given the relationships between $\succeq_{PHO^*}$ and $\succeq_{PHO}$, we can deduce from Proposition 2 that, according to $\succeq_{PHO^*}$, DA^2 performs weakly better than BOS^2 in terms of stability. The reason is that Proposition 2 shows that, according to $\succeq_{PHO}$, DA^2 performs strictly better than BOS^2 in terms of stability, and the preconditions for weak comparisons are the same for both partial orders. Also, we can deduce from Proposition 2 that, according to $\succeq_{PHO^*}$, BOS^2 does not perform weakly better than DA^2 in terms of stability. The reason is that the precondition for a strict comparison according to $\succeq_{PHO}$ precludes a reversed weak comparison according to $\succeq_{PHO^*}$. These two implications are recorded in Corollary 1.

Corollary 1. *Let the solution concept C be undominated strategy profiles. Let X denote the stable assignments correspondence. Letting F^{DA^2} and F^{BOS^2} be the functions respectively associated to DA^2 and BOS^2 by C and X, we have $F^{DA^2} \succeq_{PHO^*} F^{BOS^2}$ and $F^{BOS^2} \not\succeq_{PHO^*} F^{DA^2}$.*

Proof. **Part 1.** $F^{DA^2} \succeq_{PHO^*} F^{BOS^2}$.

By definition of $\succeq_{PHO}$ and $\succeq_{PHO^*}$, this is a direct implication of $F^{DA^2} \succeq_{PHO} F^{BOS^2}$ (Proposition 2).

Part 2. $F^{BOS^2} \not\succeq_{PHO^*} F^{DA^2}$.

By definition of $\succeq_{PHO}$ and $\succeq_{PHO^*}$, this is a direct implication of $F^{DA^2} \succ_{PHO} F^{BOS^2}$ (Proposition 2). More precisely, this follows from the fact that there exists a type profile y (given in Table 1.1) for which all undominated strategy profiles of DA^2 leads to a stable assignment whereas it is not the case of all undominated strategy profiles of BOS^2. ■

1.6 Concluding remark

The strength of the comparison between two mechanisms depends on the partial order used. One can be more confident that a mechanism will perform better than another when they can be ranked when using $\succeq_{PROP}$ than when this can only be done when using $\succeq_{PHO}$ or $\succeq_{PHO^*}$. However, given its two limitations, affirmative comparisons obtained with $\succeq_{PROP}$ are bound to be scarce and hard to obtain. In their absence, affirmative comparisons obtained with $\succeq_{PHO}$ or $\succeq_{PHO^*}$ may provide interesting indications about the respective performance to expect from two alternative mechanisms.

Appendix

1.A The school choice model

The model and notation are inspired from Haeringer and Klijn (2009). There are three students i_1, i_2 and i_3 and three schools s_1, s_2 and s_3, each endowed with one seat. Each student can be assigned to at most one school. Students have preferences over the schools they could be assigned to as well as the possibility of remaining unassigned (i.e., being self-matched). Each school has a strict priority ordering over the students. In this setting, a **(school choice) problem** is a pair $\pi = (R, \trianglerighteq)$ where

1. $R \coloneqq (R_{i_1}, R_{i_2}, R_{i_3})$ is the (strict) **preference profile** of students over the three schools, and
2. $\trianglerighteq \coloneqq (\trianglerighteq_{s_1}, \trianglerighteq_{s_2}, \trianglerighteq_{s_3})$ is the (strict) **priority profile** of schools over the three students.

The preference R_i of student i is a linear order over $S \cup \{i\}$. If student i strictly prefers school s over school s', we write $s \; P_i \; s'$. As usual, $s \; R_i \; s'$ denotes a weak preference, allowing for $s = s'$. We say that a school s is **acceptable** for a student i if $s \; P_i \; i$ and **unacceptable** if $i \; P_i \; s$. To avoid trivialities, we assume that all students find at least one school acceptable.

The priority $\trianglerighteq_s$ of school s is a linear order over the three students. If student i has a higher priority than student j at school s, then $i \vartriangleright_s j$ and we say that i is a **competitor** of j at school s. School s is a **top-priority** school for student i if

i has no competitor at school s. We denote by Π the domain of problems satisfying these assumptions.

An **assignment** is a function $\mu : \{i_1, i_2, i_3\} \to \{s_1, s_2, s_3\} \cup \{i_1, i_2, i_3\}$ that matches every student with a school or with herself. We say that student i is **assigned** in the former case, and **unassigned** in the latter case. An assignment is **feasible** if no two students are assigned to the same school.

Given any problem π, an assignment μ is **stable** if it satisfies each of the three following properties.

Individual rationality: For any student i, we have $\mu(i)\ R_i\ i$.

Non-wastefulness: For any student i and any school s, if $s\ P_i\ \mu(i)$, then $\#\{j \in I \mid \mu(j) = s\} = 1$.

No justified-envy: For any two students i and j, if $\mu(j)\ P_i\ \mu(i)$, then j is a competitor of i at school $\mu(j)$.

A **(school choice) mechanism** M is a function that associates every problem π in some domain $\Pi^M \subseteq \Pi$ of problems with a feasible assignment. We say that a mechanism is individually rational, non-wasteful or stable, if $M(\pi)$ is individually rational, non-wasteful or stable for all $\pi \in \Pi^M$. As is common, when there is no ambiguity about $\trianglerighteq$, we often use $M(R)$ to denote the assignment selected by mechanism M.

We assume that the three schools report their priority ordering truthfully to the mechanism. A **type profile** y is a school choice problem $\pi = (R, \trianglerighteq)$ (and thus $Y = \Pi$), and the **players** of mechanism M are the three students. For the two mechanisms that we consider, the **strategy space** S_i of each student i consists in the set of reported preference Q_i for which at least one school is unacceptable and at least one school is acceptable.

For any type profile y, the pair (M, y) defines a strategic form game for which students report a preference and the outcome is the assignment selected by M under the profile of reported preferences. Given (M, y), the strategy-space of student i is the set of all the preferences of i that are featured in at least one problem of Π^M. We call these strategies **reported preferences**. A **reported profile** is a list $Q := (Q_{i_1}, Q_{i_2}, Q_{i_3})$ of the reported preferences of all students.

The outcome of the game when students report Q is assignment $M(Q)$. Student i evaluates this assignment according to her true preference R_i. In particular, strategy Q_i is a **(weakly) dominant strategy** for student i if

$$M_i(Q_i, Q_{-i}) \ R_i \ M_i(Q'_i, Q_{-i}), \qquad \text{for any } Q_{-i} \text{ and any } Q'_i.$$

In turn, strategy Q_i is a **dominated strategy** for student i if

$$M_i(Q'_i, Q_{-i}) \ R_i \ M_i(Q_i, Q_{-i}), \qquad \text{for any } Q_{-i} \text{ and some } Q'_i$$

and $M_i(Q'_i, Q'_{-i}) \ P_i \ M_i(Q_i, Q'_{-i})$ for some Q'_{-i}. A strategy is **undominated** if it is not dominated.

1.B Two mechanisms

In this section we describe the two school choice mechanisms we compare, which are members of the class considered in (Haeringer and Klijn (2009)). We first describe BOS^2, a constrained version of the Boston mechanism for which students are allowed to report preferences on two schools only.

Input : A (reported) school choice profile.

Round 1: Students apply to the school they reported as their favorite school. Every school that receives more applications than its capacity starts rejecting the lowest applicant in its priority ranking, up to the point where it meets its capacity. All other applicants are *definitively accepted* at the schools they applied to, and capacities are adjusted accordingly.

Round 2 : Students who are not yet assigned apply to the school they reported as their second favorite school. Every school that receives more *new* applications in round 2 than its *remaining* capacity starts rejecting the lowest *new* applicants in its priority ranking, up to the point where it meets its capacity. All other applicants are definitively accepted at the schools they applied to. The algorithm terminates and all students not yet assigned remain unassigned.

We now turn to DA^2, a constrained version of the Deferred Acceptance mechanism for which students are allowed to report preferences on two schools only.

Input : A (reported) school choice profile.

Round 1: Students apply to the school they reported as their favorite school. Every school that receives more applications than its capacity *definitively rejects* the lowest applicant in its priority ranking, up to the point where it meets its capacity. All other applicants are *temporarily* accepted at the schools they applied to (this means they could be rejected at a later point).

Round 2 **:** Students who were rejected in round 1 apply to the school they reported as their second favorite school. Every school considers the new applicants of round 2 *together with* the students it temporarily accepted. If needed, each school *definitely rejects* the lowest students in its priority ranking, up to the point where it meets its capacity. The algorithm terminates and all students not yet assigned remain unassigned.

1.C Preliminary results on undominated strategies under DA^2 and BOS^2

Propositions 1 and 2 require identifying undominated strategies under DA^2 and BOS^2. The following lemmas provide the necessary results for such identification. They are direct implications of characterization results taken from Haeringer and Klijn (2009) and Decerf and Van der Linden (2018a).

Lemma 2. *If student i finds only one school acceptable, then reporting only this school is a dominant strategy under both BOS^2 and DA^2.*

Proof. This is a straightforward implication of the characterization of dominant strategies in constrained BOS and constrained DA in Decerf and Van der Linden (2018b). ■

Lemma 3. *Assume that the most-preferred school of student i is a top-priority school for i. Under both BOS^2 and DA^2, (1)*

i has a dominant strategy and (2) i is assigned to her most-preferred school when she plays her dominant strategy.

Proof. This is a straightforward implication of the characterization of undominated strategies in constrained BOS and dominant strategies in constrained DA in Decerf and Van der Linden (2018a) (Proposition 2 and 4). ■

Lemma 4. *Assume that the second most-preferred school of student i is a top-priority school for i. Student i has a dominant strategy under DA^2, which consists in reporting these two schools truthfully.*

Proof. This is a straightforward implication of the characterization of dominant strategies in constrained DA in Decerf and Van der Linden (2018a) (Proposition 2). ■

Let student i's **most-preferred top-priority school** be the school that i prefers among the schools that are top-priority for i (if any).

Lemma 5. *Assume that the most-preferred school of student i is not a top-priority school for i. Strategy Q_i is undominated under DA^2 only if Q_i reports two schools, Q_i ranks these two schools according to i's true preference and i weakly prefers these two schools over her most-preferred top-priority school.*

Proof. Haeringer and Klijn (2009) (Proposition 4.2) show that a necessary condition for Q_i to be undominated under DA^2 is that Q_i reports two schools and Q_i ranks these two schools according to i's true preference. Decerf and Van der Linden (2018a) (Proposition 3) show that another necessary condition for Q_i to be undominated under DA^2 is that i weakly prefers these two schools over her most-preferred top-priority school. ■

Lemma 6. *Strategy Q_i is undominated under BOS^2 if and only if (i) the school reported first is i's most-preferred top-priority school or (ii) the school reported first is not top-priority for i and Q_i reports two schools, one of which is strictly preferred to i's most-preferred top-priority school.*

Proof. This is a straightforward implication of the characterization of undominated strategies in constrained BOS in Decerf and Van der Linden (2018a) (Proposition 4). ■

Chapter 2

Sex-Selective Abortions and Instrumental Births as the two faces of the Stopping Rule. New measures and world evidence

Joint with Jean-Marie Baland and Guilhem Cassan

Abstract: The stopping rule refers to a behaviour by which parents continue child bearing till they reach a specific number of children of a given gender (boys, in general). Under this behaviour, parents can choose to carry out these pregnancies to term and raise a larger number of children than originally desired. Some of these children are therefore not desired for their own sake, and can be defined as 'instrumental'. When additional births become too costly, parents can also resort to sex-selective abortion by terminating pregnancies of the undesired gender. We argue that these two practices are the two complementary expressions of the stopping rule and ought to be considered under a unified framework.

In this paper, we take the child as the unit of interest and propose new measures of detection of these two practices. With instrumental births, a girl is, on average, exposed to a larger number of younger siblings than a boy. Under sex-selective abortion, a boy has on average more sisters among her elder siblings than a girl. These measures are easily implementable, precise, and do not rely on a natural sex ratio. We

carry out our detection tests over a large set of countries and quantify, for the countries identified by our tests, the magnitude of gender bias in parental preferences. We highlight, in particular, the minor role played by sex-selective abortion as compared to instrumental births in fertility behaviour.

2.1 Introduction

Strong preference for sons is prevalent in many societies. Many cultural factors account for these gender biased preferences, including patrilocality (Ebenstein (2014)), old age support (Ebenstein and Leung (2010); Lambert and Rossi (2016)), or the burden of the dowry (Arnold et al. (1998), see also Williamson (1976), Das Gupta et al. (2003) or Jayachandran (2015) for detailed reviews). In this paper, we explore the demographic consequences of these preferences.

Gender biased preferences lead to two fertility practices: the *"stopping rule"*, by which parents continue having children until they reach their desired number of boys[1] and *"sex-selective abortion"*, by which parents abort foetuses of the undesired gender. Both practices are viewed as distinct consequences of son preference and tend therefore to be investigated separately by the literature (Arnold (1985); Basu and de Jong (2010); Bhalotra and Cochrane (2010); Jayachandran (2015)). Figure 2.1.1 presents the number of papers in Jstor in which the words "stopping rule" alone, "sex-selective abortion" alone, and both combined appear at least once.[2] A vast majority of

[1]The preferred gender is boy in most cases, so, for simplicity, we refer to boys only, whereas these practices could also be used to target a desired number of girls.

[2]As "stopping rule" is also a term widely used in computer sciences, we have added the word "fertility" to all those searches. There exist also several synonyms to "stopping rule", which we have included in our searches The searches made were: "stopping rule" OR "differential fertility behavior" OR "son-targeting fertility behavior" OR " son-preferring fertility behavior" NOT "sex-selective abortion" ; "sex-selective abortion" NOT "stopping rule" NOT "differential fertility behavior" NOT "son-targeting fertility behavior" NOT "son-preferring fertility behavior" ; "sex-selective

papers mention only one practice, with no reference to the other. The stopping rule has also become of marginal interest as compared to sex-selective abortion. A plausible reason for this evolution is the belief that sex-selective abortion has led to the disappearance of the stopping rule. As we will show below, this belief is, to a large extent, wrong. Another possible reason is that the consequences of sex-selective abortion are thought to be more problematic than those of the stopping rule, which is highly debatable. A third reason is methodological, as sex-selective abortion distorts the observed sex ratios, making their use unfit for detecting the stopping rule.

In this paper, we argue that sex-selective abortion is not fundamentally distinct from the stopping rule. When focussing on pregnancies, sex-selective abortion and the stopping rule are essentially equivalent. Figure 2.1.2 describes the pattern of pregnancies for families which desire one son and no daughter, with or without sex-selective abortion. In both cases, their behaviour is identical: they carry on having pregnancies when the foetus is a girl and they stop having pregnancies when the foetus is a boy.[3]

In the following, *"stopping rule"* is used to refer to the general practice of childbearing until the desired gender composition is reached. This practice is made of two components: *"instrumental birth"* and *"sex-selective abortion"*. *"Instrumental*

abortion" AND "stopping rule" ; "sex-selective abortion" AND "differential fertility behavior" ; "sex-selective abortion" AND "son-targeting fertility behavior" ; "sex-selective abortion" AND "son-preferring fertility behavior" . All articles found in duplicates were removed. These searches were conducted on January 9, 2023. Non relevant results referring to agricultural practices were manually cleaned.

[3]Obviously, girls are born in one case and aborted in the other.

Figure 2.1.1: Number of articles referring to gender biased fertility practices.

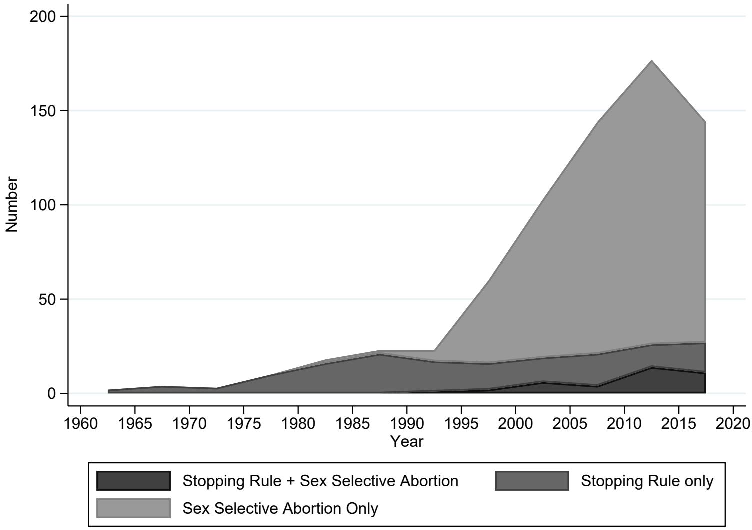

Data source: Jstor and author's computations.
Reading: in 2000-5, out of 103 articles published mentionning the words "stopping rule" (and its synonyms) or "sex-selective abortion", 84 mentionned sex-selective abortion alone, 13 mentionned "stopping rule" alone and 6 mentionned both.

birth" describes this behaviour by which parents have children until the desired gender composition is reached.[4] Under this

[4]This is the practice the literature generally refers to when using the

Figure 2.1.2: Sex-selective abortion as a stopping rule: an illustration.

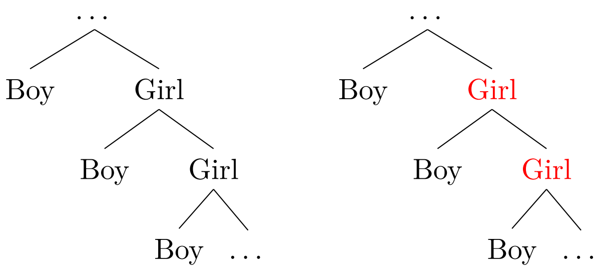

(a) Pregnancies without sex-selective abortion

(b) Pregnancies with sex-selective abortion

behaviour, children of the undesired gender are born. They are instrumental since their birth happened only as a result of their parents' attempt to get a child of the desired gender.[5] *"Sex-selective abortion"* refers to the practice of terminating pregnancies until reaching the desired gender composition. To our knowledge, the stopping rule has never been considered under that light.

We analyze these two components under a unified framework: under the stopping rule, sex-selective abortion and in-

---

term "stopping rule".

[5]Some authors refer to these children as undesired children. We believe that, ex-ante, parents practicing instrumental births know that they require these births in order to attain their desired gender composition. Therefore, these births are better termed 'instrumental' than 'undesired'. Ex-post, of course, these children may be undesired.

strumental birth are the two technologies households can use to reach their desired gender composition. When the sex-selective technology is not available, parents carry on having pregnancies and children. When sex-selection technology is available and cheap enough, parents carry on having pregnancies but terminate some of them. When costs are intermediate, sex-selective abortion does not fully replace instrumental births and the two technologies are used within the same family. In this case, parents resort to instrumental births for their first pregnancies while switching to sex-selective abortions for the last ones. Once having chosen sex-selective abortion, parents never switch back to instrumental births.

Given this equivalence, studying these two technologies jointly is important because of the policy trade-offs involved. They are indeed associated with undesirable, but different, outcomes. On the one hand, instrumental births lead to higher than desired fertility (Sheps (1963); Park (1978, 1983); Arnold (1985); Clark (2000); Dahl (2008); Basu and de Jong (2010)). It is the source of negative outcomes at the level of the society (fertility is higher than desired), the mother (for instance, through increased maternal mortality (Milazzo (2018))) and the girl (by exacerbating sibling competition, reducing birth intervals or via other forms of differential treatment (Arnold et al. (1998); Jensen (2003); Bhalotra and van Soest (2008); Jayachandran and Kuziemko (2011); Rosenblum (2013); Rossi and Rouanet (2015); Altindag (2016); Jayachandran and Pande (2017)). On the other hand, sex-selective abortion controls fertility but leads to missing girls at birth (Sen (1990); Anderson and Ray (2010); Bhalotra and Cochrane (2010); Anukriti et al. (2022)): too few girls are born as compared to boys. Abortions are costly for mothers and involve highly skewed sex ratios at the society level (Tuljapurkar et al. (1995); Hesketh and Xing (2006); Bhaskar (2011); Edlund et al. (2013); Grosjean and Khattar (2018)). By contrast, girls, once born, are more likely to be desired and face better outcomes than under the stopping rule (Goodkind (1996); Lin et al. (2014); Hu and Schlosser (2015); Kalsi (2015); Anukriti et al. (2022)).

Unless gender biased preferences are changed, a policy which increases the cost of sex-selective abortion (Nandi and Deolalikar (2013)) leads to more instrumental births. Without entering into an 'optimal population' debate, it is not clear that, given the implications for mothers, children and the society, this particular policy is desirable per se. While changing the relative costs of these technologies may not necessarily be feasible or desirable (as discussed in Das Gupta (2019), see also Mohapatra (2013)), we emphasize here their substituability, whereby a policy targeting one technology has direct consequences on the prevalence of the other. Sex-selective abortion is forbidden in many countries (Darnovsky (2009)), rising the cost of sex-selective abortion compared to instrumental births, and inducing parents to turn to instrumental births. By contrast, some countries impose strict limits on the number of births, such as the one child policy in China, increasing the realtive cost of instrumental births and inducing parents to turn to sex-selective abortions. This policy trade-off is absent from many debates about these practices.

Under our unified framework, we derive tests allowing to detect instrumental births and sex-selective abortions and quantify their relative prevalence. Our detection tests do not rely on the benchmark provided by a natural sex ratio.[6] They can be applied without any further assumption or prior knowledge about gender preferences. They are also robust to the distortions in the sex ratio induced by sex-selective abortion, as well as to underlying differences in the natural sex ratio at birth.[7] As we elaborate below, this makes our tests more reliable than existing alternatives.[8] As stressed above, these tests are important from a policy perspective as they allow the detection of both instrumental births and sex-selective abortions, and therefore inform policy makers, researchers and activists about where to target their efforts, both geographically or in terms of the technology used. In addition, our quantification exercise shows that instrumental births remain largely prevalent, even in contexts in which sex-selective abortions are widely practiced. Thus, in Haryana, an Indian state known for the high prevalence of sex-selective abortion, instrumental births are twice as prevalent as sex-selective abortion (out of the 22% of children directly affected by the stopping rule, 15 percentage points are born via instrumental births while the remaining 7 percentage points are missing due to sex-selective abortion).

The main intuitions behind these tests are as follows. Were data on pregnancies and the gender of foetuses available to

[6]See the discussion in Anderson and Ray (2010) on the absence of a universal natural sex ratio.

[7]See the discussion in Anderson and Ray (2010) on how Sub Saharan Africa countries seem to have a different natural sex ratio at birth and how this may lead to very large underestimation of the phenomena of missing women at the world level if not properly taken into account.

[8]Bongaarts (2013)'s international study of sex ratio at birth and of sex ratio of the last born for example can not distinguish between changes in natural sex ratio across space and time and changes in sex ratio caused by sex-selective abortion.

the researcher, detecting and measuring the stopping rule, instrumental births and sex-selective abortions could be done jointly. However, because data on pregnancies are typicaly not available, measuring instrumental births and sex-selective abortions has to be done through two separate, though related, approaches. They are defined at the child level and are both based on the information available in the demographic composition of the siblings.

Building on Yamaguchi (1989), Arnold et al. (1998) and Ray (1998), we show that, under the stopping rule and a preference for sons, female foetuses are on average followed by more pregnancies than male foetuses, but they have the same number and gender distribution of previous pregnancies foetuses. In the absence of sex-selective abortion, this translates into girls having on average more younger siblings than boys, but having the same number and gender distribution of elder siblings. Our formalization, while taking the perspective of the child rather than that of a family with a completed fertility, replicates some well-known consequences of instrumental births: total fertility is higher than desired (Sheps (1963)), the total number of siblings is higher for girls than for boys (Yamaguchi (1989); Basu and de Jong (2010)) and, within families, the average birth order of girls is lower than for boys (Basu and de Jong (2010)). As a matter of fact, all these are the direct consequences of girls having more younger siblings than boys under the stopping rule. Our child level approach makes formalization much simpler as well as more precise than previous attempts, and therefore contributes to the large theoretical literature on instrumental births.[9]

[9]For example the influential formalization of Basu and de Jong (2010)

This result also suggests a simple method to identify countries (or societies in general) in which instrumental births prevail, by detecting countries in which girls have more younger siblings than boys. Compared to other methods such as the sex ratio of the last born (Jayachandran (2015)), there is no need to refer to a natural sex ratio at birth, which has been shown to vary across time and space (Chahnazarian (1988); Waldron (1998); Bruckner et al. (2010)) and cannot provide a reliable benchmark as discussed in Anderson and Ray (2010). This property also makes our test robust to the practice of sex-selective abortion, which directly affects sex ratios. Our method can also be applied to families which have not completed their fertility and thereby allow us to consider recent, instead of past, behaviours (Haughton and Haughton (1998)). As we will show, besides countries in South Asia and Northern Africa, many Central Asian and European countries do implement a stopping rule.

Over the recent years, the practice of sex-selective abortion developed rapidly as a method to control the gender composition in the family. As abortions are typically not observable, the literature focusses on the evolution of the sex ratio at birth over time and birth ranks (Park and Cho (1995); Arnold et al. (2002); Hesketh and Xing (2006); Jha et al. (2006); Almond

concludes that instrumental births lead to two distinct consequences: that girls have more siblings (as in in Yamaguchi (1989)) and that girls have a higher birth order within family. They also write that "girls will be born into relatively larger family." Our approach shows that there are not two but only one consequence of instrumental births: girls have more younger siblings (but the same number of older siblings). As a consequence, girls are not born in family larger than boys: at birth, their family is exactly as large as those of boys. But their families will grow larger after their birth. Only our child level approach can make that point formally.

and Edlund (2008); Abrevaya (2009); Bhalotra and Cochrane (2010); Jha et al. (2011); Chen et al. (2013); Lin et al. (2014); Anukriti et al. (2022)). This literature exploits the empirical fact that sex-selective abortion tends to be practiced at later birth ranks: it looks at how the sex ratio at birth changes across ranks before and after the arrival of sex-selection technology. In doing so, this literature implicitely relies on the fact that, under son preference, sex-selective abortion is more likely the larger the proportion of girls among elder siblings.[10] We develop a formalization of this insight, showing that parents always prefer to postpone sex-selective abortions and turn to them when unsatisfied with the gender composition of their first births. When instrumental births and sex-selective abortions are not too costly, parents will use both: they will start with instrumental births and switch to sex-selective abortions in later births. This is because the opportunity cost of instrumental births discontinuously rises in the last births, when the birth of an instrumental girl would prevent the family from reaching its desired number of sons. This suggests a simple test of detection based on the proportion of girls among elder siblings. Absent sex-selective abortion, the proportion of boys and girls among elder siblings should be independent of the gender of the child. When sex-selective abortion is widespread, the gender composition of elder siblings differs: a girl is more likely to be born (aborted) when parents are (not) satisfied with the gender composition of her elders, that is, if she has a large proportion of boys (girls) among her elder siblings. As a result a difference in the proportion of boys among elder sib-

[10]Figure 2 in Anukriti et al. (2022) provides a nice graphical illustration of this pattern.

lings across genders is evidence of sex-selective abortion. As in the detection of instrumental births, this test does not rely on the use of a natural sex ratio.[11] Countries that practice sex-selective abortion are essentially located in South and Central Asia and Eastern Europe.

We then calibrate a simple model of gender-biased desired fertility to decompose the stopping rule into instrumental births and sex-selective abortions. This approach provides a measure of the proportion of 'instrumental' children which, as we show, is large and biased against girls. Given the general decline in desired fertility, it also tends to increase over time. In the process, we compute a desired sex ratio, defined as the ratio of the desired number of boys to that of girls, to assess gender biased preferences. Following Anderson and Ray (2010), we also provide a measure of 'missing' girls at birth and show that instrumental births remain the most prevalent stopping rule technology even in the presence of widely available sex-selection methods. Among countries practicing the stopping rule, we show that focusing on sex-selective abortion alone leads to an under estimation of the consequences of the stopping rule by more than 66%.

The structure of the paper is as follows: we first present our formalization of instrumental births and of sex-selective abortion. Following the detection methods developed in that section, we then identify the countries in which the instrumental births or sex-selective abortion prevail, and quantify their

[11]Apart from the difficulty linked to the absence of a benckmark natural sex ratio, the use of the sex ratio as birth to detect sex-selective abortion has been criticized (Dubuc and Sivia (2018)) for offering a potentially biased view of the *proportion* of parents willing to use it to reach their desired number of sons, in the presence of decreasing desired fertility.

prevalence. We discuss that type of data needed for our tests to be used and show that while fertility history data is ideal, our tests do reasonnably well with household roster data even in the presence of patrilocality and gender difference in age of marriage. Finally, we compare the result of our test to those obtained under more traditional measures and discuss these alternative approaches.

2.2 Demographic consequences of the stopping rule

2.2.1 The stopping rule and the number of younger siblings

In its analysis of the consequences of instrumental births, the theoretical literature in demography has extensively focussed on outcomes at the family level, such as the total fertility or the sex ratios among children (e.g Sheps (1963); Yamaguchi (1989); Clark (2000); Basu and de Jong (2010)). Our approach differs by taking the perspective of an individual child and by considering the stopping rule at the level of pregnancies rather than actual births. This perspective drastically simplifies the modeling effort and delivers more precise empirical predictions. It also highlights the equivalence between the two technologies behind the stopping rule.

The main intuition of our measure goes as follows: suppose that the only reason for which parents have pregnancies is to reach a desired number of boys. Each pregnancy is considered as a draw in a lottery in which having a male foetus is a “success”, while having a female foetus is a “failure”. When a male is “drawn”, parents are one unit closer to their objective. When a female is “drawn”, parents make no progress, and additional pregnancies (draws) are required in order to compensate for this failed attempt - no matter whether the foetus is sex-selectively aborted or not. This is true at each pregnancy. Therefore, compared to a male, a female foetus is a failed lottery draw, which does not contribute to the desired number of boys. As a result, a female foetus of a particular rank will be followed by exactly the same number of pregnancies as a male foetus of the same rank plus the expected number of additional draws required to have the male foetus that she is not.

We now formally investigate this mechanism. Consider first the case under which couples can have an unlimited number of children and want to have a given number b^* of boys. At any pregnancy, parents have p chances to have a boy and $(1-p)$ chances to have a girl. As a result, in a 'large' population and at each pregnancy, there is exactly $\frac{1-p}{p}$ female foetus for each male foetus. In other words, the (male to female) sex ratio at any rank in this population is constant and equal to $\frac{p}{1-p}$. The 'stopping rule' has no effect on the sex ratio in the aggregate or at each rank (Sheps (1963)). (Of course, by its very definition, the stopping rule determines the gender of the last birth and therefore the sex ratio of the last pregnancy.)

By definition, the mothers of a male or of a female foetus of pregnancy k had the same number of $k-1$ pregnancies in the past. The gender composition of these past pregnancies is also identical, as it is independent of the gender of the k^{th} foetus itself. For instance, if we assume that parents want at least 3 boys and focus on a foetus at rank 3, there are four possible combinations for the previous pregnancies: ($female-female$, $female-male$, $male-female$, $male-male$). These events occur with probability $\left((1-p)^2,(1-p)p,p(1-p),p^2\right)$, which is the distribution faced by a foetus at rank 3, independently of whether it is a female or a male. As a result, the only difference between male and female foetuses of the same rank comes from subsequent pregnancies. The critical difference between the two is the fact that having a male foetus at a particular rank implies that parents are one unit closer to their desired number of boys. Having had a female, parents are not closer to their target and, therefore, need additional pregnancies to compensate. Therefore, at any given rank, a female foetus is expected to be followed by more pregnancies than a male. Let us now assume, more realistically, that the number of children born in a family cannot exceed a given maximum, $\bar{N}$, with $\bar{N} > b^*$. This additional constraint implies that, absent sex-selective abortion, some families will not reach their desired number of boys. Consider a foetus of rank k and of gender $i = b, g$ who has e older brothers, with $e + 1 \leq b^*$ (the last inequality indicates that, at rank $k-1$, the family has not yet reached her desired number of boys, b^*). We denote by $E\left(Y_i\left(k,e\right)\right)$ the expected number of pregnancies that follow a foetus of rank k^{th}. Following the discussion above, we obtain:

Proposition 1: For any number of elder brothers e, with $e+1 \leq b^*$, the expected number of future pregnancies is strictly larger for a female than for a male foetus at any rank k, with $k < \bar{N}$:

$$E\left(Y_g\left(k,e\right)\right) > E\left(Y_b\left(k,e\right)\right),\ \forall k < \bar{N}, e+1 \leq b^*.$$

Moreover,

$$E\left(Y_g\left(k,e\right)\right)-E\left(Y_b\left(k,e\right)\right)=1/p,\ \forall e+1\leq b^*,\ \text{ when }\ \bar{N}\rightarrow\infty.$$

Proof: See Appendix 2.A

The proposition remains true, independently of the technology available to parents. In particular, when sex-selective technology is not available, this implies that, at a given birth rank, girls have, in expected terms, more younger siblings than boys. As the proposition holds for each rank, we also have, by summing over all ranks, that girls on average have a larger expected number of younger siblings. This result easily extend to a situation under which parents also desire a given number of daughters $g^* > 0$. A preference for boys in this situation simply requires that the desired sex ratio, b^*/g^*, is larger than $p/(1-p)$. Under this condition, girls still have more younger siblings. To the best of our knowledge, no other plausible mechanism can produce such an outcome.

It is easy to show that the difference in the expected number of subsequent pregnancies is increasing in $\bar{N}$. More precisely, for a given number of desired boys, b^*, the difference at any rank $k < \bar{N}$ is monotonically increasing in $\bar{N}$, as does the male to female ratio of the last pregnancy. Conversely, for a given $\bar{N}$, the difference is monotonically decreasing in the number of desired boys, b^*. Moreover, when $\bar{N}$ is very large, the difference takes a very simple expression. Having had a female instead of a male foetus, parents need one more boy in the future to compensate and therefore require, in expected terms, $1/p$ more pregnancies. As a result, at any given rank, the mother of a female foetus is expected to have $1/p$

additional pregnancies. In particular, if $p = 1/2$, she will, on average, have 2 more pregnancies than if the foetus had been a boy.[12] [13]

Figure 2.2.1 illustrates the model's main prediction with Indian data, India being a country in which the stopping rule is considered as pervasive. For all children who are at least 10 years old at the time of the survey, we have computed, at each age between -2 and 10, the average number of ever-born siblings for boys and girls separately. Before being born (at age -2 to 0), Indian boys and girls have the same number of (ever-born) older siblings. It is only after their births that the number of siblings for a girl becomes higher than for a boy, the more so the older she gets[14].

[12]These results closely parallel those of Yamaguchi (1989), who investigated the impact of the stopping rule on the expected proportion of boys in a family and total family size. We show here that these outcomes can only be driven by younger siblings.

[13]Note that a similar result can be found using an alternative model in which parents have lexicographic preferences over the number of children and the number of sons. Interestingly, a lexicographic model of the stopping rule demonstrates that the fertility squeeze (Guilmoto (2009); Jayachandran (2017)) not only applies to sex-selective abortions but also to instrumental births. See Appendix 2.B for more details.

[14]Note that if a child is not yet born (negative ages), she can only have elder siblings but when she is born (positive ages), her siblings will be either older or younger siblings. Therefore, the fact that the divergence in the number of ever born siblings emerges only after the child is born reflects that the divergence is driven *only* by younger siblings.

Figure 2.2.1: Number of ever-born siblings by age and gender in India

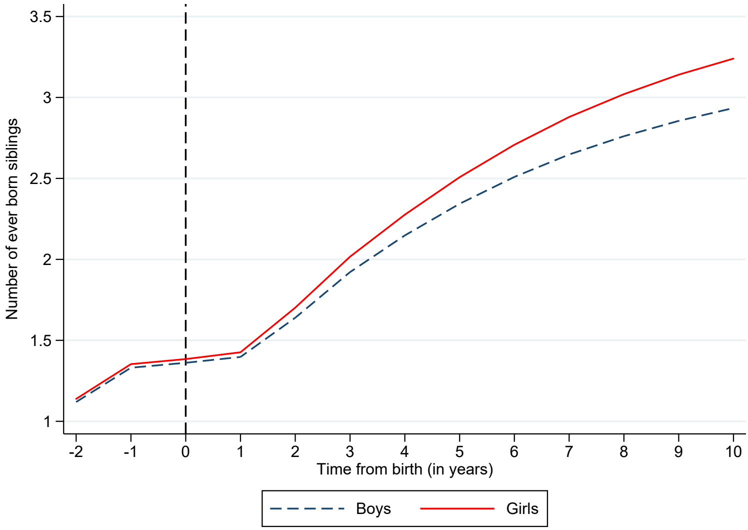

Data source: DHS India 1993, 1999, 2006, 2015, all children aged 10+ at the time of the survey.
Reading: at age 10, the average Indian girl has 3.24 ever-born siblings and the average Indian boy has 2.93 ever-born siblings.

2.2.2 Self-selective abortion and the composition of elder siblings

Detecting instrumental births is relatively straightforward, given the availability of data on birth history. A different approach is needed to detect sex-selective abortion since reliable information on pregnancies and abortions is not available. Suppose again that parents can have a maximum of $\bar{N}$ children. To simplify the discussion, we assume here that parents want exactly b^* boys and no girls, with $\bar{N} > b^*$. As discussed later, the argument easily extends to the more general setting in which $g^* > 0$. The 'natural' probability of having a boy out of each pregnancy is given by p. Each abortion implies a cost to parents of C_{ssa}, and girls have no value per se. In this simple framework, parents will always delay abortion as long as this is feasible, i.e., as long as the number of possible births left allows them to achieve their objective of b^* boys. Indeed, abortion at pregnancy j implies a cost of C_{ssa} while, by postponing to the next rank, this cost only occurs with probability $(1-p)$. In other words, future abortions are, in expected terms, always less costly than the current one. This implies that, once a couple decides to practice sex-selective abortion, other abortions are also carried out in the future in the event of other female foetuses. We therefore have:

Lemma: Once, for a given pregnancy, sex-selective abortion is chosen, sex-selective abortion is chosen for all future pregnancies.

In families which only had daughters in their previous pregnancies, the result above implies that, since parents want b^* boys, sex-selective abortion starts to be applied at birth rank $\bar{N} - b^* + 1$. More generally, for those families that did yet not achieve their desired number of boys, female foetuses are systematically aborted at all pregnancies for which the number of births left available corresponds to the number of boys that are still missing to reach their objective. In particular, full sex-selection is expected in the last rank, $k = \bar{N}$. In other words, at each rank $k \geq \bar{N} - b^* + 1$, sex-selective abortion is applied by some families to obtain some of the boys born at that rank.

That sex-selective abortion is applied at later ranks is supported by a large body of evidence (see e.g. Lin et al. (2014)). Figure 2.2.2 illustrates this idea in the case of Armenia. We report in the Figure the average sex ratio for all births of a particular rank, before and after the widespread utilization of the ultra-sound technology (2000). While before 2000, their sex ratio does not vary across ranks, the picture after 2000 is particularly striking, with a steep increase in the observed ratio in the later ranks. Thus, the sex ratio at rank 4 reaches 164, implying a proportion of boys in all births of rank 4 of 62.1%. By contrast, the sex ratio at rank 1 is identical to that prevailing before 2000, indicating a negligible amount of sex-selective abortion at that rank.

An important implication from the above argument is that sex-selective abortions are practiced in families with not enough sons and too many daughters. This observation implies that, for all boys of rank $k \geq \bar{N} - b^* + 1$ the proportion of girls among their elder siblings is larger for a boy than for a girl of the same rank.[15]

Proposition 2: Under self-selective abortion, at any rank $k \geq \bar{N} - b^* + 1$, the proportion of girls among elder siblings is larger for a boy than for a girl.

[15] Note that the probability that the sibling of rank $k - 1$ is a girl is also larger than for a girl of the same rank. Also, the proportion of girls among younger siblings is lower for girls than for boys of a given rank.

Figure 2.2.2: Sex ratio at birth by birth rank in Armenia before and after 2000

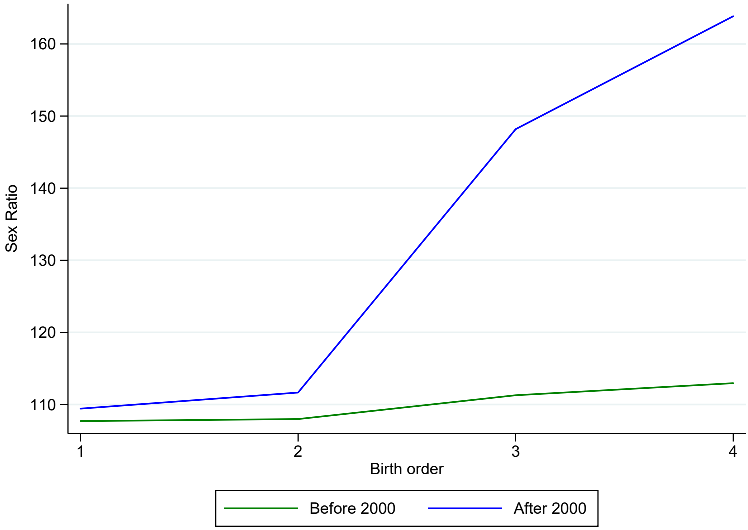

Data source: DHS Armenia 2000, 2005, 2010 and 2016, all children born at the time of the survey.
Reading: Prior to 2000, the sex ratio at birth at rank 3 is 111. After 2000, it is 148.

The proposition is at the heart of our test to detect the occurrence of sex-selective abortion. The argument easily extends to a setting in which parents also want a given number of daughters. As a matter of fact, even if gender preferences are biased in favour of girls, so that $b^*/g^* < p/(1-p)$, the proportion of boys among elder siblings is larger for a girl than for a boy, so that the proportion of girls among elder siblings is again larger for a boy than for a girl. As a result, the proposition simply states the consequences of biased preferences, regardless of the preferred gender. Finally, in the particular case in which $b^* + g^* = \bar{N}$, sex-selective abortion is already practiced in the first rank.

2.2.3 Combining Costly Instrumental Births and Sex-Selective Abortion

As argued above, instrumental births and sex-selective abortions are two complementary mechanisms driving the demographic composition of families. When sex-selective abortion becomes available, parents choose the best practice by comparing the cost of abortion to the cost of an (additional) instrumental child. We consider the case in which $\bar{N} > b^* + g^*$ and, for simplicity, we assume that all costs are constant. As above, C_{ssa} stands for the cost of abortion, while the cost of an instrumental birth is given by C_{ib}. We also have to consider the additional cost of failing to reach the desired gender composition, which we denote by C_f. This cost occurs when the number of births left available is lower or equal to the number of boys or girls that still remain to be born. In other words, the cost of instrumental births changes discontinuously when it comes at the cost of not reaching the desired gender composition: the first instrumental births are relatively cheap draws, the last instrumental births are expensive draws, while sex-selective abortions are a costly way to cheat the lotery.

Comparing these costs, three cases are possible. In the first case, $C_{ssa} > C_{ib} + C_f$ at all ranks, and abortions, being too costly, are never practiced. When $C_{ssa} < C_{ib}$ at all ranks, abortion is cheaper than an additional instrumental child, and parents always resort to sex-selective abortions for each foetus of the undesired gender. More interesting is the case in which the cost of an abortion is smaller than the combined costs of an instrumental birth and of not reaching the desired gender composition for some ranks, $C_{ib}+C_f > C_{ssa}$, but greater than the cost of an instrumental birth at lower ranks: $C_{ssa} > C_{ib}$. In this case, parents choose instrumental births for the first pregnancies and switch to sex-selective abortions at later ones, when the number of births left available is just equal to the number of desired boys or girls still required.

As the cost of sex-selective abortion falls, its prevalence increases as parents turn away from carrying out to term their last pregnancies. These cases are illustrated in Figure 2.2.3, where $\tilde{N}$ represents the number of 'cheap draws'.[16] When abortion costs are small and lower than C_{ib} at all ranks, sex-selective abortion is practiced from the first pregnancy on, all born children are of the desired gender and no instrumental births occur (Figure 2.2.3c). By contrast, when abortion costs are very large, sex-selective abortion is never practiced and, at each rank, part of the births are "instrumental" (Figure 2.2.3a). Policies affecting the number of sex-selective abortions will therefore result in opposite changes in the number of instrumental births. This framework also accomodates poli-

[16]Formally, $\tilde{N} = \underbrace{(\bar{N} - k + 1)}_{Remaining\ draws} - \underbrace{(max(0, b^* - b) + max(0, g^* - g))}_{Remaining\ required\ successes}$, with b and g the number of boys and girls already obtained at that rank and b^* and g^* the desired number of boys and girls.

cies controlling the number of births allowed to parents. For example, the "One child policy" in China makes $\tilde{N} \leq 0$.

Figure 2.2.3: Technology choice as a function of relative costs and 'cheap' draws

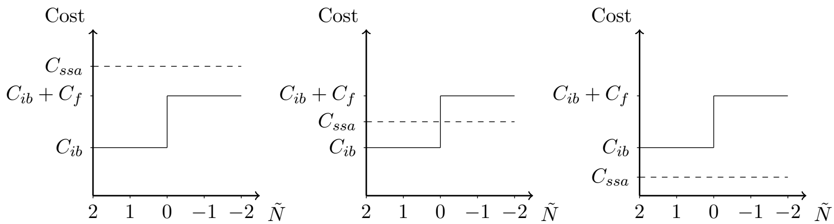

(a) Only Instrumental Births

(b) Instrumental Births followed by Sex-selective Abortions

(c) Only Sex-selective abortions

Sex-selective abortions do not therefore imply the disappearance of instrumental births. In the case illustrated in Figure 2.2.3b, parents in early ranks prefer to carry out the pregnancies while turning to sex-selective abortions in later ranks when these become necessary to reach their desired target. To illustrate this case, consider a family which desires two boys and one girl, with $p = 1/2$ and $\bar{N} = 5$. Among all possible family compositions, suppose that we observe the following sequence : girl, girl, boy, girl and finally a boy. Among these children, the girl born at rank 2 is 'instrumental' since she was not desired per se, but is born as the result of the parental desire to have two boys. Similarly, the girl at rank 4 is also instrumental. Finally, at the last rank, parents, having had three girls and one boy, will abort in the event of a female foetus. As a result, the boy born at the last rank is either the result of a natural birth, with probability p, or of an abortion with probability $1-p$.[17] In other words, $1-p$ girls should have been born in the last rank, have been replaced by a boy and should be considered as 'missing'. Among these five births, we therefore have two 'instrumental' and $1-p$ 'missing' children.

Two measures of interest can be defined here. The first one is the *share of instrumental births*, defined as the number of instrumental births divided by the total number of children, and which we interpret as the probability that a child taken at random is 'instrumental'. The second is the *share of missing children*, defined as the number of births that would have been of a different gender in the absence of abortion, again divided by the total number of pregnancies. It corresponds to the probability that a random child is born as the result of at least one previous abortion.

It is worth noting that, under the assumptions made above, the sum of these two measures is exactly equal to the share of

[17]Strictly speaking, in this case, multiple abortions are possible in the event of a sequence of female foetuses (Dimri et al. (2019)). We therefore implicitly assume that parents can have a large number of pregnancies, even though the maximal number of children is given. Under this assumption, we can infer the corresponding expected number of abortions necessary to obtain the boy who replaces the 'missing' girl at rank 5 as $1/(1-p)$. As a result, the expected number of abortions in the last rank is exactly equal to 1 ($= (1-p)*1/(1-p)$). If the number of possible pregnancies is limited, the expected number of abortions lies between $(1-p)$ and 1. Note that, given the Bernoulli process assumed, this expected number quickly converges to 1, even for a limited number of pregnancies. Thus, if at most six pregnancies in the last rank are possible, the corresponding expected number of abortions is larger than 0.98.

instrumental births that would have occurred in the absence of abortion. This is because our measure of missing children does not (and cannot) measure the actual number of abortions, which remain unobserved, but, instead, the number of foetuses ('potential' instrumental children) that have been replaced by a child of the desired gender. Absent abortion, these foetuses would be born and counted as instrumental births. For instance, in our simple numerical example, consider the set of families of five children starting with a sequence (g, g, b, g) for the first four children. In the absence of abortion, a proportion $(1 - p)$ of families present the sequence (g, g, b, g, g) and a proportion p of families, (g, g, b, g, b). On average, therefore, $2 + (1 - p)$ children are instrumental. When abortion is available, the last female foetus is replaced by a male in $(1 - p)$ families, and the only sequence observed is (g, g, b, g, b), with two instrumental children and $(1 - p)$ missing girl, which corresponds exactly to the $(1 - p)$ instrumental child above. This property strikingly illustrates this idea that instrumental births and sex-selective abortion are the two complementary dimensions of the stopping rule.[18]

Even when sex-selective abortions substitute for instrumental births, the latter remain quantitatively significant. To illustrate this point, we computed the share of instrumental births and missing children observed in families that desire 1 girl and either 1 or 3 boys, with a probability $p = 0.5$. Figure 2.2.4 presents the evolution of the two measures for different levels of the maximum number of children (up to 8).

[18] This property holds as long as the availability of abortion does not change the actual number of births, nor the desired number of boys and girls. Our assumptions of given preferences and of a given maximum family size satisfy these two requirements.

Figure 2.2.4: Decomposition of stopping rule between instrumental births and missing births

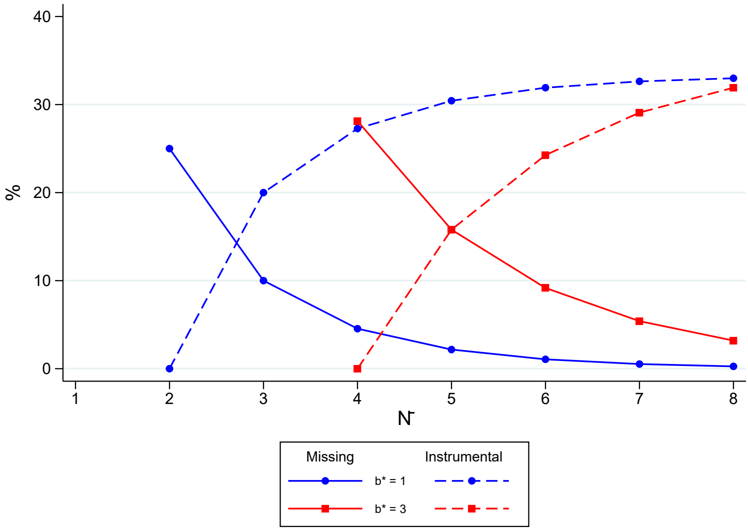

Data Source: Author's simulations.
Reading: For a desired number of boys b^* of 1, of girls g^* of 1 and a maximum number of births $\bar{N}$ of 3, there are on average 30% of children born under the stopping rule: 20% because of instrumental births and 10% because of sex-selective abortions.

Even if abortion is available, the share of instrumental births remains sizeable and, in general, exceeds the share of missing children. Thus, when parents desire 3 boys and 1 girl, with a maximum family size of 6, the share of instrumental births is equal to 24.3% (out of which 22.5% are instrumental girls and 1.8% instrumental boys) while the share of missing births is equal to 9.2% (out of which 8.9% are missing girls and 0.3% missing boys). In the absence of abortion, one would have observed 33.5% instrumental children. Also, the share of 'missing' children becomes quickly negligible when the family size is large. For instance, focusing again on the case in which parents desire 3 boys and 1 girl, the share of missing children falls down to 3.2% when the maximum family size is equal to 8 (as compared to a share of instrumental births equal to 31.2%). It is only when the number of desired boys and girls is very close to the maximum number of children that the share of missing births gets relatively large. Thus, when parents can have up to 5 children, the share of missing and the share of instrumental children are both equal to 15.8%. By construction, when the desired number of boys and girls is exactly equal to the maximum family size, no instrumental births are observed and the share of missing children is equal to 28.1%.[19]

2.3 Prevalence of the stopping rule across countries

Given the demographic consequences of the stopping rule, our theory offers a precise and straightforward strategy to detect

[19] This is in line with (Jayachandran (2017)) who highlights the increased occurrence of abortions under declining fertility. In the context of our model, we indeed observe an increase in the share of missing children, and therefore in the occurrence of abortion, as the maximum family size decreases.

and measure the prevalence of the stopping rule, without relying on partial evidence or priors about the prevailing practices. We first derive detection tests based on the sibling composition of each child. For instrumental births, our test measures the number of younger siblings she has while, for sex-selective abortion, we focus on the gender composition of her elder siblings. We present our main results at the world scale, before proposing measures of the relative importance of instrumental and missing children in countries that implement the stopping rule. We finally provide a more detailed analysis of the Indian case.

In our empirical analysis, we use the Demographic and Health Survey (DHS) of all countries available. This represents 82 countries, 2,995,509 mothers and 10,361,884 births. The DHS are particularly valuable to us as they are comparable across countries, and record the fertility history of ever married women aged 13 to 49. We can therefore reconstruct for each child at any age the number of siblings, older or younger, she had. Appendix 2.D lists the countries and surveys we used, as well as the corresponding number of observations. For the detection of sex-selective abortion, we focus on children born after 2000, since the ultra-sound technology was not widespread enough before that date. This restricts our sample to 1,685,160 mothers and 3,754,614 births in 69 countries.[20]

[20]Ideally, we would have liked to also analyze the cases of Korea, Japan, Taiwan or China. Unfortunately, appropriate data were either not available or not directly comparable to the information provided by the DHS.

2.3.1 Detecting the stopping rule

Instrumental Births

As shown in the preceding section, the difference in the number of younger siblings between girls and boys of any rank is exactly zero in the absence of stopping rule. This benchmark does not depend on a natural sex ratio, nor on whether a particular child is a last born or not. Our measure of instrumental births simply compares the difference in the number of younger siblings of girls and boys to this logical benchmark. When instrumental births prevail in favour of boys, this difference is necessarily greater than zero. When favoring girls, this difference is smaller than zero. As discussed in Section 2.2, it can also be aggregated across children and families in a straightforward manner.

To illustrate our approach, we run the following estimations for India and Bolivia:

$$\text{Nb Younger Siblings}_{it} = \sum_{t=0}^{T}(\alpha_t * \text{age}_{it} + \beta_t * \text{female}_i * \text{age}_{it}) + \epsilon_{it} \tag{2.1}$$

$$\text{Nb Younger Siblings}_{ik} = \sum_{k=1}^{K}(\alpha_k * \text{rank}_{ik} + \beta_k * \text{female}_i * \text{rank}_{ik}) + \epsilon_{ik} \tag{2.2}$$

where Nb Younger Siblings$_{it}$ is the number of ever born younger siblings of child i at age t and Nb Younger Siblings$_{ik}$ is the number of ever born younger siblings of child i at rank k.

For each child between zero and ten year old or between rank 1 and 6, we record all her younger siblings at each age, with an age or rank-varying number of younger siblings. (We also cluster standard errors at the primary sampling unit and weight each observation by the DHS sample weight.) As our estimates describe parental preferences, there is no *a priori* reasons to include additional controls in the specifications. Figure 2.3.1 below reports our estimates.

Figure 2.3.1 strikingly illustrates the prevalence of instrumental births in India, as the number of younger siblings at all relevant ages or ranks is systematically larger for girls than for boys. (These results replicate, in a regression format, the descriptive statistics presented in Figure 2.2.1 above.) As expected, this differential increases with age, to reach an average of 0.3 extra siblings at age 10. By contrast, for ranks, the relation is non-monotonic as the number of additional children declines at higher ranks. We also report the corresponding estimates for Bolivia, for which no such differential exists. At any rank, at any age, the average Bolivian girl has the same number of younger siblings as the average Bolivian boy.

We now provide the test for the prevalence of instrumental births across all the countries surveyed. Since the difference in the number of younger siblings prevails at all ranks and ages, we use equation 2.3 to estimate for each country a condensed version of equations 2.1 and 2.2[21]:

[21] Note that, when aggregating over all ranks, the average birth order of girls is lower than that of boys (Basu and de Jong (2010)). Not controlling for ranks therefore makes our measure higher (in absolute value). Since the 'aggregate-rank' effect is a direct consequence of instrumental births, and as our measure solely aims at detecting the prevalence of instrumental births, this is not an issue.

Figure 2.3.1: Differential number of ever-born younger siblings by age and rank, India and Bolivia

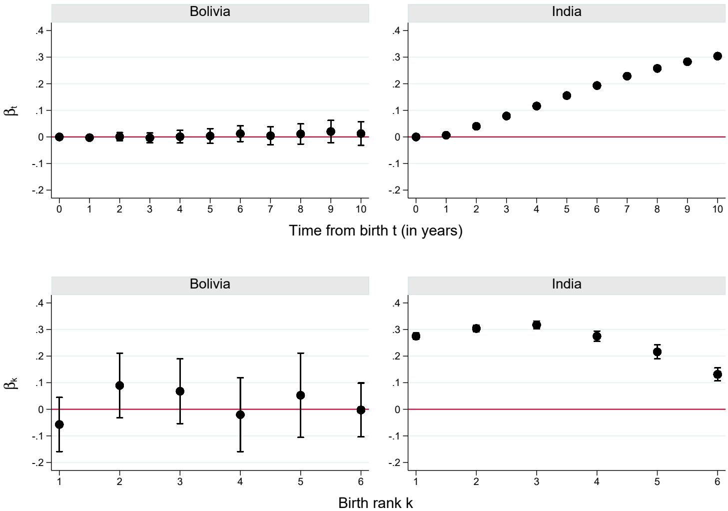

Data Source: DHS Bolivia 1989, 1994, 1998, 2003 and 2008 and DHS India 1993, 1999, 2006 and 2015.
Reading: In India, at age 10, girls have on average 0.3 more younger siblings than boys of the same age. Girls born at rank 5 have on average 0.24 more younger siblings than boys of the same rank. No difference across gender is perceivable either by age or by rank in Bolivia.

$$\text{Nb Younger Siblings}_i = \beta * \text{female}_i + \epsilon_i \tag{2.3}$$

The coefficient β corresponds to the difference in the average number of younger siblings a girl faces compared to a boy. Two additional remarks are in order. First, as illustrated by Figure 2.3.1, one could use children of a particular age to carry out our test. When measuring the difference in the number of younger siblings at younger ages, the measure is increasingly influenced by birth spacing which may vary across gender (see in particular Jayachandran (2015)). This in itself is not a issue for a measure of detection of the instrumental births, as shorter birth spacing associated with the less desirable gender simply translates into a larger number of younger siblings at a young age. On the other hand, focussing on older children implies that our measure applies only to a more distant past and does not provide information on more recent years. Given these two trade-offs, we choose here to focus on all children, but our main results are robust when focussing on children of a specific age. Second, one could also choose a particular rank over which to apply our measure and focus, for instance, on the eldest child of each family. In theory, a test on the difference in the number of younger siblings between first-born girls and boys provides a necessary and sufficient condition for the detection of the instrumental births. This is because the gender composition of younger siblings is given by a probability distribution so that a test on the eldest child carries enough information for the test to apply and requires only to know the gender of the eldest child and the number of younger siblings. In doing so however, we neglect useful information related to the consequences of the gender of siblings of higher ranks. In a 'small' sample, focussing on a particular rank therefore provides a sufficient but not a necessary test of the instrumental births.

Figure 2.3.2: Differential number of ever-born younger siblings of girls, by country

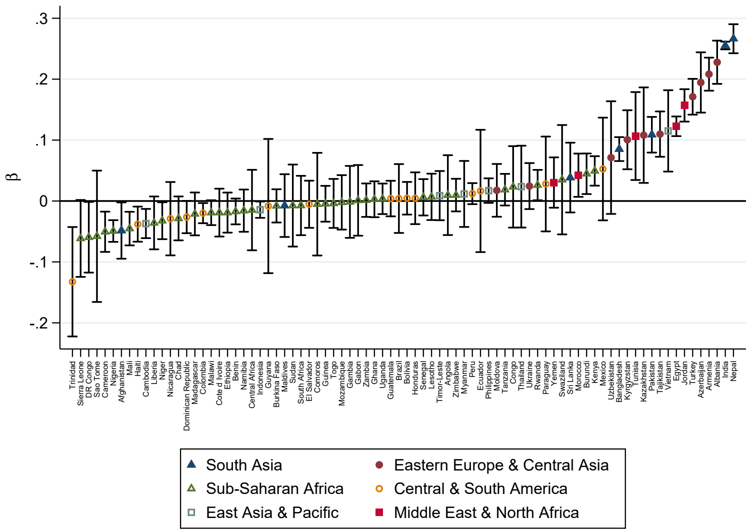

Data Source: All DHS.
Reading: In Nepal, girls have on average 0.27 more younger siblings than boys.

We report in Figure 2.3.2 the differential number in younger siblings of girls for all countries present in our sample, by increasing order. On the right-hand side of the Figure, one finds a substantial cluster of countries with a very high difference in the number of younger siblings, indicating the prevalence of instrumental births in these countries. The latter does not only include the 'usual suspects', such as Nepal, India, Pakistan or Bangladesh, but also countries of Eastern Europe and Asia, such as Albania, Turkey, Armenia, Azerbaijan, Jordan, Kazakhstan, Kyrgyzstan, Tajikistan or Vietnam, and Northern Africa, such as Egypt, Morocco or Tunisia. This fact has been essentially ignored by the economic literature (Ebenstein (2014) is a notable exception). Second, a few countries (Cambodia, Cameroon, Colombia, DR Congo, Haiti, Indonesia, Mali, Niger, Nigeria and Trinidad), display a much smaller (in absolute value) but negative coefficient, suggesting the presence of a stopping rule favoring girls and not boys.[22] This possibility is hardly mentioned in the economic literature (Williamson (1976)), but our gender neutral approach allows to identify such a case. Most countries from Sub-Saharan Africa do not apply the stopping rule. (This last statement has to be qualified, however, owing to the relative prevalence of polygamy in most of these countries. This may have implications that we discuss in the subsection 2.3.2.)

Sex-selective Abortion

Our measure of sex-selective abortion compares at the child level the gender composition of his or her elder siblings and does not require a particular benchmark, such as the one given by a natural sex ratio.[23] When sex-selective abortion does not

[22]This pattern is grossly consistent with countries in which brideprice- rather than dowry - is practiced, such as in Indonesia (Ashraf et al. (2020)) or Sub Saharan African countries (Corno et al. (2020))

[23]An alternative measure could simply focus on the gender of the preceding child, but this strategy does not use the full information available to

apply, the gender distribution among elder siblings is identical across boys and girls of any rank so that no difference can emerge. By contrast, when sex-selective abortion applies, boys tend to have more sisters and girls more brothers among their elder siblings.[24] The prevalence of sex-selective abortion can be illustrated using the following estimation:

$$\text{Sh elder girls}_{it} = \sum_t (\gamma_t * \text{year}_{it} + \delta_t * \text{male}_i * \text{year}_{it}) + \epsilon_{it} \quad (2.4)$$

where Sh elder girls_{it} is the share of sisters among alive elder siblings of a child i born in year t.[25] Note that, unlike our detection test for instrumental births, the coefficient of interest does not vary with age. This is because the composition of the elder siblings is given and does not vary with the age of the child. By contrast, we carry out our estimations at different birth years t so as to compare the current situation to that prevailing before the spread of ultra-sound technologies. Alternatively, we can also estimate the prevalence of sex-selective

the parents at the time of pregnancy. While less efficient, it also provides a sufficient condition for detecting sex-selective abortion.

[24]Note that it is also the case that boys have more sisters among their younger siblings compared to girls. Nevertheless, even when sex-selective abortion is not applied, the gender distribution among younger siblings is not the same for boys and girls, as girls have more younger siblings under the stopping rule. This may become a problem in small samples.

[25]Given that the decision to abort selectively depends on the composition of the family at the time of the pregnancy, we focus on elder siblings alive at the time of birth.

abortion for children at specific ranks k over a given period:

$$\text{Sh elder girls}_{ik} = \sum_k (\gamma_k * \text{rank}_{ik} + \delta_k * \text{male}_i * \text{rank}_{ik}) + \epsilon_{ik} \quad (2.5)$$

Focussing again on India and Bolivia, the top panel of Figure 2.3.3 below presents our estimates over several birth cohorts starting in the mid-seventies. On the bottom panel, we report the corresponding estimates for children of different ranks born after 2000.

As in the analysis of instrumental births, India and Bolivia offer a contrasting image. While Bolivia appears essentially gender neutral, with no noticeable differences between girls and boys, India exhibits a strong prevalence of sex-selective abortion for all ranks once the ultra-sound technology became widely available at the end of the nineties. Since then, the incidence of sex-selective abortion increases monotonically.

We now provide a test of sex-selective abortion across all countries which, for the sake of presentation, is based on a simplified version of equation 2.5, averaging over all ranks:

$$\text{Sh elder girls}_i = \delta * textmale_i + \epsilon_i \quad (2.6)$$

The δ coefficient, estimated separately for each country, corresponds to the difference after 2000 in the average proportion of girls among elder siblings of girls as compared to boys. Figure 2.3.4 presents these estimates by increasing order of magnitude.

On the right-hand side of the Figure, we find the countries with a significant difference in the gender composition of elder siblings. As expected, these are less numerous than in the detection of instrumental births, owing to the limited availability

Figure 2.3.3: Differential share of girls in elder siblings by period and rank, India and Bolivia

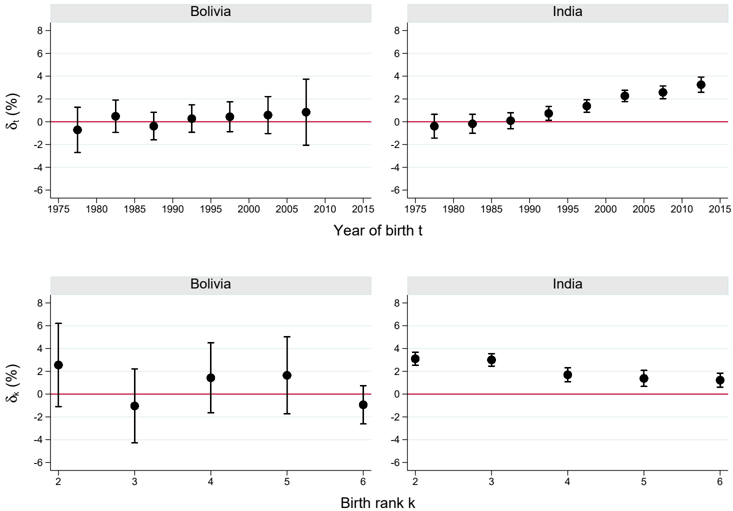

Data Source: DHS Bolivia 1989, 1994, 1998, 2003 and 2008 and DHS India 1993, 1999, 2006 and 2015.
Reading: In India, from birth cohort 1990 onwards, boys start having a larger proportion of girls among their elder siblings than girls. There is no such difference in Bolivia. For births taking place after 2000 in India, at each birth rank, boys have a larger proportion of girls among their elder siblings than girls. There is no such difference in Bolivia.

Figure 2.3.4: Differential share of girls in elder siblings of boys, all countries

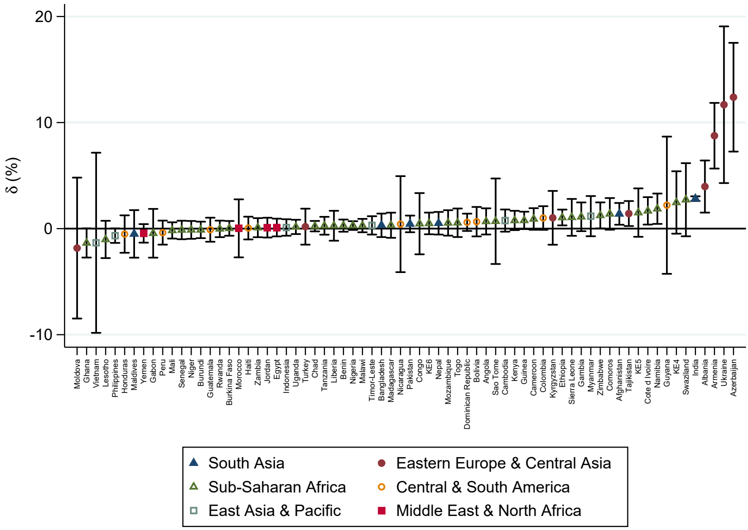

Data Source: All DHS, births taking place after 2000.
Reading: In Azerbaidjan, boys have on average a proportion of girls among their elder siblings larger by 12.39 percentage points as compared to that of girls.

of the ultra-sound technology. Among the countries identified, one finds India, Albania, Armenia, Azerbaidjan and Tajikistan in which the stopping rule also prevails, but also Ukraine. At much lower levels of significance, one also finds Cote d'Ivoire, Ethiopia, Namibia and Zimbabwe. As discussed in Section

2.3.1, our measure remains silent as to whether sex-selective abortion favors a particular gender. The observed sex ratio is needed to infer the prevailing gender bias.

2.3.2 Limits of our approach

We now discuss more systematically these two tests. The measures we propose are by themselves meaningful. Thus, our test of instrumental births directly provides a measure of the number of additional siblings and, therefore, the sibling competition a girl is exposed to compared to a boy. Our test of sex-selective abortion is a direct measure of gender diversity within families. However, in terms of gender preferences, our measures do not lend themselves to straightforward interpretations. They indeed take values that depend on the desired number of boys and girls as well as the maximal family size, which vary across time and space. As a result, our measures cannot be directly used to compare the intensity of gender preferences across countries.[26] In this respect, our measures therefore only provide sufficient conditions for the prevalence of instrumental births and sex-selective abortions.

Second, our measures only capture biases in gender preferences. For instance, if families have the same desired number of boys and girls ($\frac{b^*}{g^*} = \frac{p}{1-p}$), no differential across gender can arise and our two measures are equal to zero. As a result, a non-conclusive detection test cannot differentiate between families that have no preference regarding the gender of their children and families that apply the stopping rule to achieve an equal number of boys and girls. A similar remark holds if,

[26]As discussed in McClelland (1979), this qualification also holds for other classical measures, such as the parity progression ratio.

in the population, parents have opposite gender preferences, with about half of them applying the stopping rule in favour of boys and the other half in favour of girls. Clearly, all widely used tests also suffer from this shortcoming. What we actually detect through our tests is whether preferences are on average biased towards a particular gender in a population.

Third, our two measures should be implemented concurrently in order to assess the prevalence of the stopping rule, as they refer to two separate mechanisms of the same fundamental behaviour. Moreover, since sex-selective abortion tends to be applied at later ranks, it does not neutralize the consequences of the stopping rule in earlier ranks but makes instrumental births increasingly harder to detect empirically at later ranks. By contrast, instrumental births have, by themselves, no impact on the detection of sex-selective abortion since they cannot affect the gender composition of older siblings.

An additional difficulty comes from the possibility of a selective recall bias. Under-reporting children has two consequences on our measures: on the one hand, when computing our measures on children of rank k, some children of rank k are missing, which leads to missing observations in that rank and a possible selection bias; on the other hand, those children will not be accounted for when computing our measures on their siblings; leading to a measurement bias. As long as the recall bias is gender neutral, so that boys are as likely to be under-reported than girls, some observations are missing, but our measures remain unbiased. In demographic surveys, the main recall bias come from under-reporting elder girls who died in early age. In our measure of instrumental births, these 'forgotten girls' reduce the number of younger siblings of their elders by the same amount, independent of their gender. As a result, the difference in the number of younger siblings of these elders remains unchanged. There is therefore no measurement bias at this level. However, some elder girls, with a larger number of younger siblings, are systematically not accounted for (the selection bias). This biases downwards our measure, which still provides a sufficient condition for the prevalence of instrumental births.

Concerning sex-selective abortion, there are no clear reasons to believe the recall bias to cause a systematic selection bias in our estimates, as long as the girls unaccounted for present a gender distribution among their elder siblings that is similar to that of an average girl in the sample. For children that are observed, however, there are good reasons to believe that these forgotten girls affect differently the gender composition of elder siblings. One indeed expects underreporting to be more frequent in families with a stronger son preference. In these families, the under-reported girls tend to be followed by more younger brothers. The fall in the proportion of girls in elder siblings therefore affects the observed boys much more than the observed girls. As a result, boys on average present an even lower proportion of sisters in their elder siblings compared to girls of the same rank, leading to a downward bias of our measure. Our test again provides a sufficient condition.[27] We also discuss in Appendix 2.C the impact of other potential observational biases on our tests, as those present in household rosters which systematically omit elder children living outside

[27] By contrast, traditional measures based on the observed sex-ratio are systematically affected by this under-reporting, over-estimating the occurrence of sex-selective abortion.

the household.

Finally, our two measures are not appropriate in all settings. First, they require parents to have on average more than one child. Our measures can not be used to analyze, for instance, the one child policy in China and its demographic consequences.[28] The sex ratio at birth, with all its shortcomings, turns out to be the only measure available. Second, our test for instrumental births applies essentially to monogamous societies. In polygamous settings, one cannot exclude the possibility that men having a strong preference for boys choose to have more children with the wife that give them a son at first birth. Under this argument, mothers with a female first born have fewer children and boys, on average, end up having a larger number of younger siblings than girls. (This however points to a limitation of regular surveys which do not collect systematic information on the father of the child.)

2.3.3 Measuring the prevalence of the stopping rule

We now intend, on the basis of the previous estimates, to quantify the prevalence of the stopping rule. We want, in particular, to estimate (1) the gender bias in abortion rate and the resulting share of missing children, (2) the desired fertility by gender, and in particular, the difference between the number of desired boys and the number of desired girls and, finally, (3)

[28]This limitation also holds for all measures relying on family size to detect instrumental births (such as the parity progression ratio or the measures proposed by Basu and de Jong (2010); Yamaguchi (1989); Rossi and Rouanet (2015)) or on ranks to detect sex-selective abortion (Bhalotra and Cochrane (2010)).

the share of instrumental children by gender. As is now clear, our previous tests do not provide by themselves such indicators and we need to compute these relying, on the one hand, on a particular value of the natural sex ratio and, on the other hand, on a simple model that structures parental preferences. We assume this model to be common for the set of families under analysis (e.g. a country over a given period). After describing our empirical strategy, we first present the estimated temporal evolution of some critical indicators for India. We then provide a summary table of these indicators for all the countries in which a gender bias was detected through our two tests, before returning to the Indian case, which we investigate in more details by state and caste.

Empirical Approach

We first measure sex-selective abortion by comparing a natural sex ratio to the observed sex ratio, along the spirit of the methodology proposed by Anderson and Ray (2010). Note that this simple comparison does not provide an estimate of the total number of abortions per gender, since these cannot be inferred using the actual number of children born by gender. We are therefore unable to estimate for each gender separately the number of abortions or 'replacements' that may have occurred. What we can do instead, through this comparison, is to estimate the excess number of 'replacements' against a particular gender as compared to the other. For the sake of presentation, we again assume that abortion rates are biased against girls, so that this comparison provides us the share of missing girls. Let N_b and N_g stand for the observed number of boys and girls. We first compute, given the number of existing children, the counterfactual population of girls which we should observe under the natural sex ratio, where p stands for the 'natural' probability of a boy at each birth. We refer to this number as the potential population of girls, N_g^P, defined as follows:

$$N_g^P = (1-p)(N_b + N_g)$$

This expression allows us to define the share of missing girls at birth among all children, m_g:

$$m_g = \frac{N_g^P - N_g}{N_b + N_g}$$

For instance, suppose that we observe in a population of 200 children 110 boys and 90 girls. With a natural sex ratio of 100 girls for 100 boys, we should have observed a potential population of 100 girls, and the proportion of missing girls is then equal to 10/200, that is 5%.[29]

In a second step, we measure the share of instrumental children, by estimating the desired number of boys and girls. In order to do so, we calibrate a simple model of a representative

[29]The proportion of missing girls computed here differs from that in Anderson and Ray (2010) since we compute the number of girls that should 'replace' boys under the natural sex ratio instead of the additional number of girls that should have been born given the number of boys observed. We therefore rely on the actual population as a natural benchmark, keeping total population fixed, while they consider a potential population of children that should be alive but are not observed. The rationale for using a measure of potential population in Anderson and Ray (2010) lies in their focus on adult excessive mortality, while our measure of instrumental children, used later, requires us to focus on children that are actually born.

household which would like to have a given number of boys, b^*, and a given number of girls, g^* for a maximum family size given by $\bar{N}$. We define X as the total number of births necessary to obtain b^* boys and g^* girls, given a probability of male birth equal to p. Under a discrete approach, the probability distribution of X is the sum of two truncated negative binomial distribution, and is given by the following expression[30]:

$$P(X = x|b^*, g^*, p) = \binom{x-1}{b^*-1} p^{b^*} (1-p)^{x-b^*}$$
$$+ \binom{x-1}{g^*-1} (1-p)^{g^*} p^{x-g^*}$$

$$\text{for} \quad x \in \mathbb{N} \in \{b^* + g^*, \bar{N}\} \quad \text{and} \quad b^*, g^* \geq 1$$

The first term of this expression represents the probability to have $b^* - 1$ boys in the first $x - 1$ births and a boy at the x^{th} birth. The second term similarly represents the probability to have $g^* - 1$ girls and b^* boys in the first $x - 1$ births, and a girl at the x^{th} birth. In the following, we rely on a continuous version of this expression, in which $b^*, g^* \in \mathbb{R}_+$ and the binomial coefficients are replaced by Gamma functions:

$$f_X(x; b^*, g^*, p) \propto \frac{\Gamma(x)}{\Gamma(b^*)\Gamma(x - b^* + 1)} p^{b^*} (1-p)^{x-b^*}$$
$$+ \frac{\Gamma(x)}{\Gamma(g^*)\Gamma(x - g^* + 1)} (1-p)^{g^*} p^{x-g^*}$$

[30]Note that when the desired number of children of one gender, g^* for instance, is equal to 0, the probability distribution of X reduces to a simple negative binomial distribution with parameters b^* and p.

$$\text{for} \;\; x > b^* + g^* \;\; \text{and} \;\; b^*, g^* \geq 0$$

Under this expression, the distribution of the number of younger siblings for boy (X_b) of the a first rank is given by $f_{X_b}(x; b^* - 1, g^*, p)$. Similarly, the distribution of the number of younger siblings for a girl (X_g) of the first rank is given by $f_{X_g}(x; b^*, g^* - 1, p)$.[31] Our empirical strategy relies on the fact that the number of younger siblings of the first born, given his (her) gender, provides all the information needed in terms of family size and composition (given the distribution above). This property is directly related to our previous observation according to which focussing on the first born is, in a large sample, necessary and sufficient for the detection of the stopping rule. We first compute μ_b and μ_g, the average number of younger siblings for first-born boys and girls observed in the sample. Given a large enough number of observations, we know that $\mu_b \to E(X_b|b^* - 1, g^*, p)$ and $\mu_g \to E(X_g|b^*, g^* - 1, p)$, the expected number of younger siblings for a first-born boy or girl given the distribution above. Given particular values of p and $\bar{N}$, we then compute the expected value of X_b and X_g for all possible values of b^* and g^*.[32]

[31] Strictly speaking, one needs to add a normalizing multiplicative constant for this expression to integrate to 1. This constant, which we will estimate, depends on b^* and g^* but quickly converges to 1 for large enough values of b^* or g^*.

[32] The value of p used in these computations corresponds to the currently observed probability of a male birth across the population and takes into account the fact that some sex-selective abortions already took place. The value of $\bar{N}$ is chosen to be equal to the 90th percentile in the number of children observed in the country under analysis. Our results are essentially unaffected by the choice of this particular value as compared to the 80th,

We then select the values of b^* and g^* which minimize the distance between the observed means μ_i and the corresponding expected values $E(X_i)$ for $i = b, g$, where the distance is defined as the sum of the differences in absolute value. To obtain the number of instrumental children of a particular gender, we simply compute the difference between the actual and the desired number of children.[33]

Measuring instrumental and missing births across countries

We start by illustrating the evolution of gender preferences in India. For that country, we compute the desired numbers of children by gender, over intervals of five years starting in 1975 (corresponding to the year of birth of the child concerned). Figure 2.3.5 below presents the estimated desired total fertility (by summing the number of desired boys and girls) on the left axis as well as the the proportion of desired boys among these, which is a direct measure of gender biased preferences (on the right axis).

Over the whole period, desired fertility decreased in India in which it fell from about 3.3 in 1980 to 1.75 in the recent years. The proportion of desired boys in total desired fertility is relatively stable at around 60%, which corresponds to a 'desired' sex ratio of 150 boys for 100 girls $(0.57/(1-0.57))$.

We then compare the desired to the actual number of girls and boys and compute the share of instrumental boys and girls.

95th or 99th percentiles.

[33] It is important to note that our approach is valid as long as sex-selective abortions are applied in the last ranks and simply replace the gender of the last born without affecting the actual number of births.

Figure 2.3.5: Desired fertility & Proportion of desired boys in India

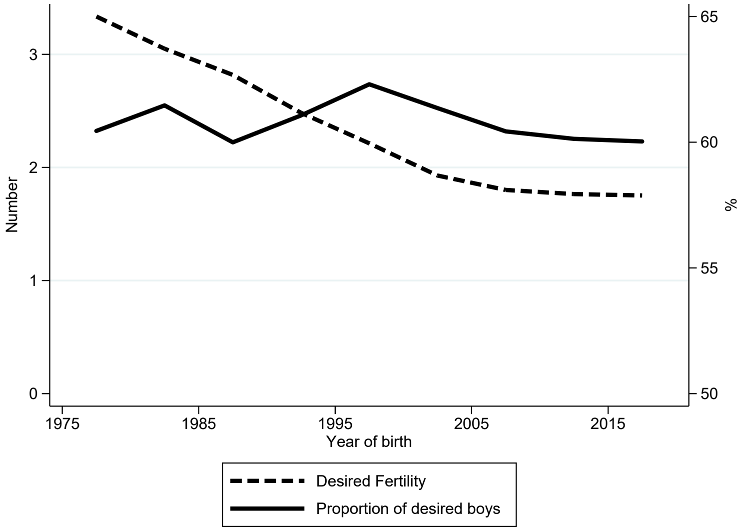

Data Source: DHS India 1993, 1999, 2006 and 2015.
Reading: Over the period 1995-2000, the total desired fertility is 2. The desired proportion of boys is 62%.

We also compute the share of missing girls at birth using each country's pre-1980 sex ratio as estimated in Chao et al. (2019) as the natural sex ratio at birth.[34] Figure 2.3.6 reports for an average family the number of desired and instrumental chil-

[34]See Appendix 2.E for a list of the ratio used.

dren, separately for boys and girls, as well as the number of missing girls. Over the whole period, the desired number of boys and girls decreased monotonically, while the number of instrumental children remains constant, which implies that the share of instrumental boys and girls increased throughout. As expected, girls are also more likely to be instrumental. Starting after 1990, the share of missing girls at birth is by contrast modest and remains stable.

We now replicate our approach to estimate the share of instrumental and missing children for all countries which we identify as applying the stopping rule. Using all births occurring after 2000, we identify 18 countries: Albania, Armenia, Bangladesh, Cameroon, Colombia, Comoros, DR Congo, Egypt, India, Jordan, Kenya, Nepal, Niger, Pakistan, Rwanda, Tajikistan, Turkey and Yemen. Table 2.3.1 reports the following indicators: the desired family size, the desired sex ratio,[35] the actual sex ratio, the proportion of instrumental boys among alive boys, the proportion of instrumental girls among alive girls, the share of instrumental children, the share of excess instrumental girls[36] and the share of missing girls at birth. The last column presents the prevalence of the stopping rule by summing the share of excess instrumental girls and that of missing girls at births. The prevalence of the stopping rule

[35] Defined as above as the ratio between the number of desired boys and the number of desired girls.

[36] Computed as the difference between the number of instrumental girls and the number of instrumental boys divided by the total number of children. Only the share of excess instrumental girls can be compared to that of missing girls at birth, since the latter is, strictly speaking, the net difference between the share of missing girls and the share of missing boys (which we cannot separately observe nor estimate given our approach).

Figure 2.3.6: Desired fertility, proportion of instrumental children and missings girls at birth in India

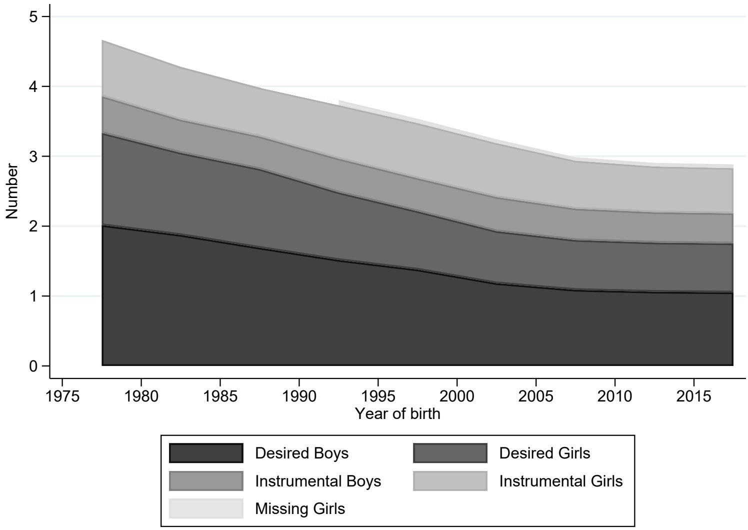

Data Source: DHS India 1993, 1999, 2006 and 2015.
Reading: Over the period 2010-2015, the number of missing girls is 0.05, that of instrumental boys 0.69 and of instrumental girls 1.02 on average per family.

is the share of children directly affected by the stopping rule, either because they are instrumental or because they are born as a result of sex-selective abortions. Most of these countries display a strong bias in preferences for boys, with a desired sex ratio that varies between 109 (for Kenya) and 232 (for

Armenia), and is particularly large in Asian countries as it never falls below 117 (in Pakistan), largely above the actual sex ratios.

Given this bias, the proportion of instrumental girls is systematically larger for girls than for boys. In Armenia, for instance, 64.5% of girls can be considered as instrumental, as against 25.3% of boys. In India, the corresponding figures are 53.4% and 34.2% for girls and boys respectively. Overall, the proportion of instrumental children hovers around 30%, with a proportion of instrumental girls close to twice as large as that of boys. On the other hand, 5 countries, Colombia, Cameroon, Comoros, Niger and DR Congo display a bias in preferences towards girls. The prevalence of the stopping rule is very diverse, with up to 21.8% children affected in Armenia and as compared to 1.6% in Kenya or 0% in Colombia.[37] Decomposing the stopping rule between its instrumental births and missing births components underlines the overall predominance of instrumental births.

The proportion of missing girls at birth, which we can only estimate for countries identified as practicing sex-selective abortion in the relevant sample, reaches 4.4% in Armenia, 1.7% in Albania, 3% in Tajikistan and 3.7% in India. The corresponding shares of excess instrumental girls are estimated at 17.4% for Armenia, 14.6% for Albania, 6.8 % in Tajikistan and 8.5% for India, largely exceeding missing girls. In other words, countries practicing the stopping rule overwhelmingly practice instrumental births only. In the countries in which both in-

[37] Our test detects Colombia as practicing the stopping rule. However the difference in the number of younger siblings between girls and boys, while significant, is small enough so that our model does not detect biased preferences.

strumental births and sex-selective abortion is practiced, the latter is at lower magnitude compared to instrumental births. For stopping rule countries taken as a whole, stopping rule affects 9.1% of children, more than two third of which (6.4%) via instrumental births. Therefore, focusing on sex-selective abortion only, as is typically done in the literature, leads to an underestimation of more than 66% of the prevalence of the stopping rule.

Table 2.3.1: Preferences & Fertility

	Desired family size	Desired sex ratio	Actual sex ratio	Instrumental boys (%)	Instrumental girls (%)	Instrumental children(%)	Excess instrumental girls (%)	Missing girls (%)	Stopping Rule (%)
Armenia	1.26	232	112	25.26	64.44	43.9	17.42	4.35	21.77
Albania	1.38	207	101	25.95	59.95	42	14.61	1.73	16.34
Tajikistan	2.54	173	106	16.46	37.13	25.44	6.81	2.98	9.79
Jordan	2.74	169	102	21.98	44.48	32.21	8.23	0	8.23
Nepal	1.81	166	107	27.43	51.07	38.58	9.59	0	9.59
Rwanda	3.06	155	106	14.06	28.22	20.23	4.36	0	4.36
Egypt	2.2	153	105	23.85	42.27	32.38	6.77	0	6.77
India	1.64	148	110	34.17	53.36	43.52	8.49	3.74	12.24
Yemen	3.37	146	110	17.45	31.06	23.58	4.41	0	4.41
Bangladesh	1.76	123	107	28.47	37.51	32.83	3.36	0	3.36
Turkey	1.62	122	109	35.69	45.2	40.35	3.99	0	3.99
Pakistan	3.1	117	105	26.83	33.34	29.98	2.32	0	2.32
Kenya	2.34	109	100	35.33	39.32	37.3	1.59	0	1.59
Colombia	1	100	104	57.62	57.62	57.62	0	0	0
Cameroon	2.76	83	109	27.61	20.72	24	-2.26	0	2.26
Comoros	2.99	82	100	32.81	24.86	28.67	-2.78	0	2.78
Niger	3.87	80	103	19.57	13.47	16.29	-1.81	0	1.81
DR Congo	3.59	77	103	25.87	17.09	21.15	-2.77	0	0
Total	1.94	138	108	32.27	47.04	39.4	6.62	2.42	9.17

We end this empirical investigation by focussing on the case of India more finely, distinguishing between states and castes. We first replicate our approach over the 17 largest states of India, for all births occuring after 2000. Table 2.3.2 reports the same indicators as above for each state, which we rank by decreasing order of the desired sex ratio. As widely documented in the literature (following Sen (1990)), we observe a strong divide in gender preferences between North-Western and Southern states. Thus, the desired sex ratio is as high as 246% in Gujarat, 243% in Haryana or 202% in Punjab, but falls down to 123% in Tamil Nadu or 129% in Andra Pradesh. (The actual sex ratio follows closely this ranking, from the exceptionnally high 125% and 119% in Haryana and Punjab to around 106% in Southern States.) Thus, in Gujarat and Haryana, for each desired girl, parents desire almost 2.5 boys. These strong biases in the desired sex ratios in the North imply a very large proportion of instrumental girls: 58.3% of girls are instrumental in Haryana, as compared to 19.1% of the boys. By contrast, in regions in which the desired sex ratio is more balanced, the share of instrumental girls and boys are much closer: in Andhra Pradesh, for instance, these two shares are much closer from one another.

Table 2.3.2: Preferences & Fertility across Indian States and Castes

	Desired family size	Desired sex ratio	Actual sex ratio	Instrumental boys (%)	Instrumental girls (%)	Instrumental children(%)	Excess instrumental girls (%)	Missing girls (%)	Stopping Rule (%)
					States				
Gujarat	1.66	246	114	20.18	60.43	38.32	16.16	4.74	20.9
Haryana	1.68	243	125	19.14	58.27	36.5	15.21	6.87	22.08
Punjab	1.36	202	119	27.36	60.63	43.23	14.62	5.65	20.27
Himachal Pradesh	1.26	200	106	31.48	64.76	47.89	15.96	2.92	18.88
Madhya Pradesh	1.84	197	106	23.88	54.85	38.17	12.45	2.94	15.39
Bihar	2.41	180	107	19.09	43.39	29.84	8.54	3.19	11.74
Maharashtra	1.42	168	114	31.04	55.94	43.05	10.91	4.77	15.68
Odisha	1.67	165	107	26.02	48.94	36.73	9.02	0	9.02
Rajasthan	1.85	164	114	29.03	52.47	40.19	9.78	4.69	14.47
Uttar Pradesh	2.72	157	112	14.08	28.67	20.42	4.51	4.11	8.62
Goa	1.03	151	107	39.11	59.49	49.27	10.05	0	10.05
Karnataka	1.25	150	106	38.26	58.24	48.18	9.64	0	9.64
West Bengal	1.02	149	108	45.22	64.63	55.12	10.81	0	10.81
Assam	1.94	140	106	23.81	37.81	30.36	5.01	0	5.01
Andhra Pradesh	1.6	129	106	25.37	35.98	30.41	3.8	2.92	6.72
Tamil Nadu	1.25	123	109	36.98	47.11	41.96	4.36	3.45	7.81
Kerala	2.07	105	105	0	0	0	0	0	0
					Castes				
High Castes	1.42	246	114	21.59	62.55	40.41	17.07	4.64	21.72
Other Backward Caste	1.83	186	111	22.53	50.14	35.1	10.55	4.07	14.62
Scheduled Tribe	1.9	175	106	26.54	52.63	38.78	10.61	2.83	13.44
Muslims	2.29	160	108	24.62	45.61	34.36	7.95	0	7.95
Scheduled Caste	1.82	156	109	30.43	51.67	40.61	8.92	3.44	12.36

In the Southern states, the practice of sex-selective abortion is undetected in Karnataka, Goa and Odisha, and remains relatively negligible in other States. By contrast, self-selective abortion is widely practiced in the Northern States of Haryana and Punjab, where the shares of missing girls are as high as 6.9% and 5.7%. The share of missing girls in the Northern States fluctuates around 4%. This being said, the excess instrumental girl share remains largely above those figures, indicating again that, in the implementation of the stopping rule, the practice of instrumental births remains preponderant. For instance, excess instrumental girls represent 15.2% of children in Haryana while the share of missing girls is order of magnitude lower, at 6.9%. That is, in Haryana, 22.1% of children are affected by the stopping rule, but focusing on sex-selective abortion alones leads to an understimation of more than two third of the prevalence of the stopping rule.

We report in the lower panel of the Table a similar exercise distinguishing between five social groups: the High Castes[38], the Other Backward Castes, the Scheduled Castes, the Scheduled Tribes and the Muslims.[39] There again, our estimates follow the established Caste hierarchy, which matches closely the observation that gender biased preferences are stronger among higher castes (Chakravarti (1993); Kapadia (1997); Field et al. (2010); Luke and Munshi (2011); Cassan and Vandewalle (2021)). Thus, the desired sex ratio is on average equal to 246% among high castes but falls to 140% among Muslims. The share of missing girls ranges from 4.7% among High Castes to

[38]Defined here as all individuals non belonging to the other categories.

[39]The Muslim category includes Muslims classified as Other Backward Classes, our Other Backward Classes category therefore only contains non Muslim individuals.

0% among Muslims, much inferior to the share of instrumental births.

2.4 Comparison with other approaches

2.4.1 The sex ratio of the last born

In the literature, the most popular measure of instrumental births is based on a literal interpretation of the stopping rule: the last born in the family tends to be a boy. As a result, countries in which the stopping rule is prevalent should display a large proportion of sons among the last born. This proportion is then simply compared to a counterfactual sex ratio, typically the “natural” sex ratio. While intuitive, this approach suffers from some important shortcomings. First and foremost, there is no universal natural sex ratio at birth: depending on the ethnic group and the period considered, it can vary between 103 and 107 boys per 100 girls (see for instance Chahnazarian (1988); Waldron (1998) and the discussion in Anderson and Ray (2010)). Sex ratio at births also vary with environmental factors, nutritional status or paternal age, so that, even within the same ethnic group at a particular time period, one cannot rely on a well-defined benchmark (see, for instance, Bruckner et al. (2010) and Catalano and Bruckner (2005)).[40] The existence of such a benchmark is made even more elusive in the

[40]For instance, Catalano et al. (2008) show that women under colder weather abort more male foetuses, so that a 1ř C increase in annual temperature predicts one more male per 1,000 females born in a year. In a similar vein, Helle et al. (2009) in their analysis of sex ratios between 1865 and 2003 showed a strong increase of excess male births during periods of exogenous stress, such as World War II.

presence of sex-selective abortion.[41]

Second, by focussing on the gender of the last born, this approach naturally applies to families with completed fertility. As a result, this measure necessarily describes the behavior of older cohorts of mothers. By contrast, through our test, the difference in the number of younger siblings emerges as soon as families have reached a number of births exceeding their desired number of boys or girls, largely before completing fertility. Moreover, this difference can be detected at each rank, so that a test carried out at the level of the eldest child is already informative (particularly in the event of sex-selective abortion at late ranks). Our measure can therefore detect behavioral changes much sooner than measures relying on the sex ratio of the last born.

We now compare the relative performance of our test to that based on the sex ratio of the last born. As explained above, the latter requires a natural sex ratio to be used as a reference. Given the uncertainty surrounding its precise value, we run the test for two plausible values of this ratio, 103 and 107.[42] In Figure 2.4.1 below, we compare the value obtained under our measure (on the vertical axis) to the corresponding value of the sex ratio of the last born (on the horizontal axis). Each dot in the graph corresponds to a particular country (the corresponding confidence intervals are not reported for the sake of exposition). The two measures are, as expected, reasonably correlated.

---

[41] Additionnally, the use of survey data makes the estimates particularly noisy: thus, for an observed sex ratio of 105, the 95 percent confidence interval ranges between 100.8 and 109.2 in a sample of 10,000 births.

[42] Clearly, the comparison is even less favourable to the sex ratio of the last born for more extreme, but plausible, values of the benchmark.

Figure 2.4.1: Our test of instrumental births against the sex ratio of the lastborn

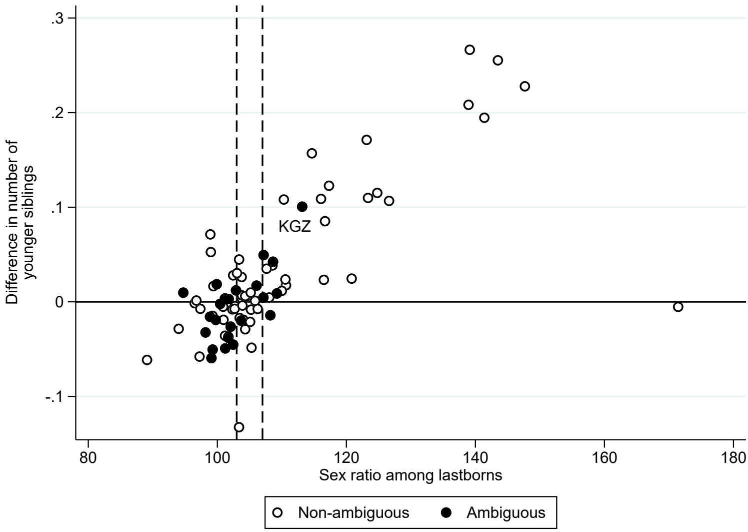

Data Source: All DHS
Reading: In Kyrgyzstan, girls have 0.1 more younger siblings than boys and the sex ratio of the lastborn is 113. However, the sex ratio of the lastborn is not statistically different from both 103 and 107 and does not allow to conclude that instrumental births are used.

The Figure illustrates the poor performance of methods based on the sex ratio of the last born. It is for 71% of the countries (the white dots) that the observed sex ratios, given their confidence intervals, lead to conclusions that do not depend on the value chosen as a reference (103 or 107). (Our test agrees with 84% of them.) By contrast, for the other countries (the black dots), the conclusion is ambiguous. We report these cases with their corresponding confidence intervals in Figure 2.4.2 below. According to our test, the stopping rule prevails in half of them (which we represent by black triangles or squares).[43] Thus, Kyrgyzstan, Morocco and Kenya, for instance, apply the stopping rule against girls according to our test, but fail to be detected by the sex ratio of the last born when a cut-off ratio of 107 is used.

2.4.2 Other popular measures

Another method used in demography is the "parity progression ratio" (Ben-Porath and Welch (1976); Williamson (1976); Arnold (1997); Arnold et al. (1998); Norling (2015)). It evaluates, at a given birth rank, the relative probability to continue childbearing (the opposite of being the last born) given the gender of the child at that rank. This measure, while close to the "sex ratio of the last born", is particularly relevant here as it does not rely on a natural sex ratio at birth. It however suffers from a number of limitations. First, it is a rank-specific measure, with no clear interpretation when the measure gives conflicting results at different ranks.[44] Relatedly, given that it

[43]As stressed above, our test provides a sufficient condition for instrumental births and may thus leave a number of situations undetected.

[44]For example, Filmer et al. (2009)'s find evidence of instrumental births as detected with parity progression for families of size 3, but not for families of sizes 2 and 4 in Sub Saharan Africa. They write "it is difficult to take in all of the coefficients at a glance." In addition, given the com-

Figure 2.4.2: Added precision of our instrumental births test compared to the sex ratio of the lastborn

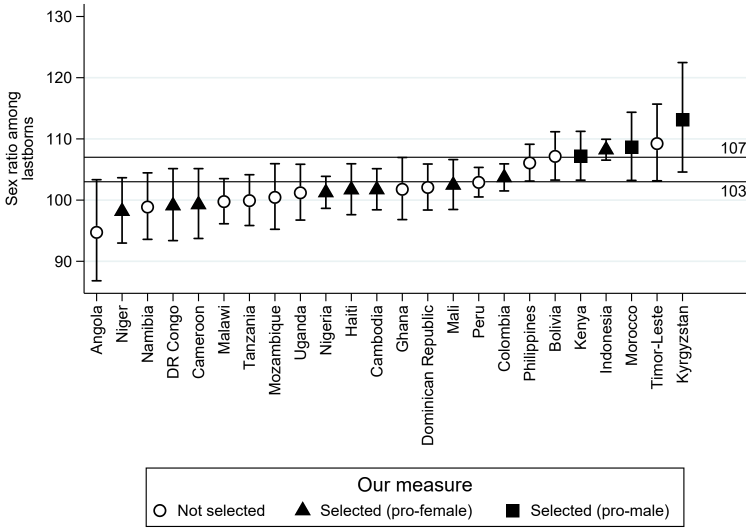

Data Source: All DHS data of countries considered as ambiguous in Figure 2.4.1.
Reading: In Kyrgizstan, the sex ratio of the last born is 113, but is not statistically different from 107. Our test allows to unambiguously classify Kyrgizstan as practicing instrumental births.

plexity of the approach, they restrict their analysis to comparing families not having any sons at given rank to families not having any daughter, therefore omitting from the analysis all intermediate cases.

is based on a ratio of two probabilities, the literature does not provide a clear way to aggregate it over ranks. One possibility could be to estimate, over all ranks, the difference (instead of the ratio) between boys and girls in their probability of having a younger sibling. This difference however corresponds essentially to the sex ratio of the last born, and uses exactly the same information. In this respect, our measure generalizes this approach by counting the number of younger siblings obtained and thereby better exploiting the information available. The parity progression ratio, because of its focus on the next pregnancy, is less efficient.

Second, children of all ranks below the 'desired number of boys or girls' necessarily have younger siblings, irrespective of their gender. Thus, if parents want, for instance, at least 2 boys, the first born of the family will necessarily have a younger sibling. It is only at later ranks that the parity progression ratio can detect a stopping rule behavior. This is problematic for comparative studies, as the desired number of sons and daughters may vary across countries and over time and would require to vary the rank analyzed across countries according to their desired number of boys and girls.

One may also think of using birth-spacing as a measure of instrumental births, following the idea that parents with a strong preference for sons will reduce the time between a newborn girl and her next sibling (see Jayachandran (2015); Rossi and Rouanet (2015)). One can then compare the average birth spacing of a girl compared to a boy, possibly aggregated over all ranks. Under this approach, the only reason why parents would want to selectively reduce birth spacing is because they want more younger siblings when the new born is a girl. However, the reduction of birth spacing is not a necessary step to do so. Therefore, while the detection of a gender difference in birth spacing implies the practice of instrumental births, the opposite is not true. Moreover, birth spacing is a particularly noisy observation, given the uncertainty associated with pregnancies. Finally, since sex-selective abortion affects birth spacing, this measure becomes less relevant when sex-selective abortion becomes widespread (Dimri et al. (2019)).

A last set of measures proposed in the literature is based on the consequences of the stopping rule. According to Basu and de Jong (2010), girls tend to be born in larger families and to have, within families, earlier ranks than boys (see also Yamaguchi (1989)). However, as shown in Section 2.2, girls do not face larger families at birth. It is only after their birth that their families grow larger under the stopping rule. This also explains why, within families, girls are born at earlier rank than boys on average. The measure we propose follows the same intuition but is more direct and precise. Following the idea behind the parity progression ratio, Arnold (1985) proposes to compare the declared use of contraceptives depending on the gender composition of the family. Being based on current use, this method is less sensitive to recall biases, but may suffer from report biases as it relies on sensitive information. It also crucially hinges upon the availability of contraceptives.

2.4.3 The sex ratio and sex-selective abortions

The theoretical literature on sex-selective abortion is less abundant. In fact, the literature mostly focused on detecting its occurrence following the introduction of the ultra-sound technology. The dominant approach rests on a comparison between the actual and the natural sex ratio in a given population. More sophisticated methods rely on the idea that sex-selective abortion is less prevalent in first ranks. This literature typically follows a difference-in-difference approach by comparing sex ratios at birth across ranks and over time, for countries such as Taiwan (Lin et al. (2014)), South Korea (Park and Cho (1995)), China (Zeng et al. (1993); Chen et al. (2013)), India (Bhalotra and Cochrane (2010); Jayachandran (2017); Anukriti et al. (2022)) or the United States (Abrevaya (2009)). Our approach of sex-selective abortion gives a theoretical foundation to these empirical studies, while allowing for a measure that is not rank-based and is, therefore, better suited for comparative approaches.

As above, we now assess the relative performance of our test to that of the traditional approach, which compares the observed proportion of boys in the population to the natural sex ratio (which, as above, we assume to be either 103 or 107). Figure 2.4.3 reports the value obtained under our measure (on the vertical axis) and the corresponding value of the observed sex ratio (on the horizontal axis).

The white dots represent countries for which a simple comparison of sex ratios leads to conclusions that do not depend on the value of the cut-off (103 or 107). They barely represent 43% of the countries (our measure agrees with 69% of these conclusions). Conversely, the black dots represent all these countries for which the choice of the benchmark leads to conflicting conclusions about the prevalence of sex-selective abortion. These ambiguous cases represent the remaining 57% of our sample. In Figure 2.4.4, we report for these countries a classical mean test using the observed sex ratio with the corre-

Figure 2.4.3: Our test of detection of sex-selective abortion against the sex ratio at birth

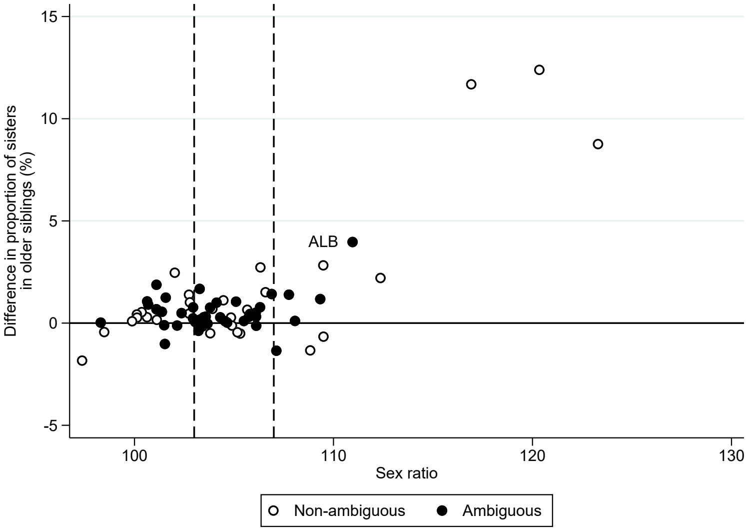

Data Source: All DHS, births taking place after 2000.
Reading: In Albania, boys have 3.97 percentage points more girls in their elder siblings than girls and the sex ratio at bith is 111. However, the sex ratio at birth is not statistically different from both 103 and 107 and does not allow to conclude that sex-selective abortions are used.

sponding 95% confidence intervals. Among these, 7 countries (18 %) are identified by our test as practicing sex-selective abortion (as indicated by the black dots). This is the case, for instance, of Albania or Afghanistan for which the traditional approach remains inconclusive with a benchmark of 107.

Figure 2.4.4: Added precision of our test of sex-selective abortion compared to the observed sex ratio

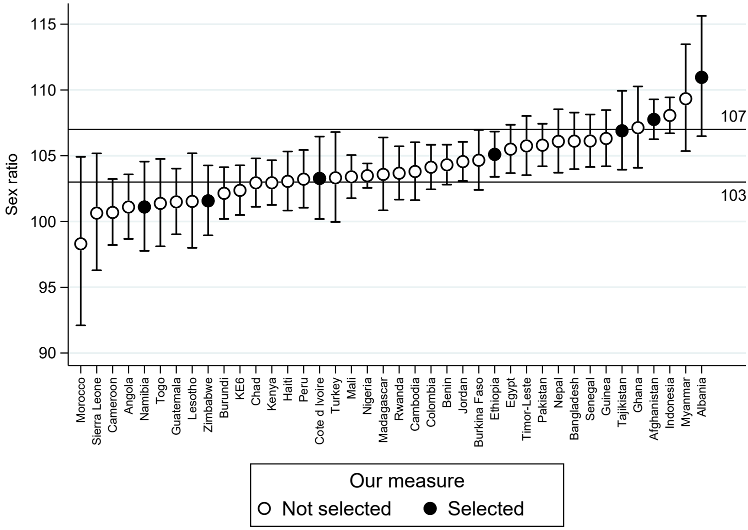

Data Source: All DHS (births taking place after 2000) of countries considered as ambiguous in Figure 2.4.3.
Reading: In Albania, the sex ratio at birth is 111, but is not statistically different from 107. Our test allows to unambiguously classify Albania as practicing sex-selective abortions.

2.5 Conclusion

The stopping rule refers to this behaviour by which parents continue child bearing until they reach a specific number of children of a given gender. Parents can then choose to carry out these pregnancies to term, leading to a larger number of children than originally desired, a practice defined as instrumental births. They can also choose to abort foetuses of a specific gender, a practice known as sex-selective abortion. While these two practices have been investigated independently in the literature, they are closely related as they both result from the same fundamental behaviour. We propose a unified framework to consider them jointly. This framework underlines the policy trade off implied by the substituability of the two practices.

Were pregnancies directly observable, these two practices could be measured in a straightforward manner. The literature provides different indirect methods aimed at estimating the consequences of these practices, which suffer from important shortcomings. Taking the child as the unit of interest, we propose, with the help of a simple model, new measures to detect these two practices. Under instrumental births, a girl is, on average, exposed to a larger number of younger siblings than a boy. Under sex-selective abortion, a girls also has on average more elder brothers than a boy. Unlike the existing measures proposed in the literature, our measures do not require the use of a counterfactual benchmark. They can be easily implemented, are defined at the level of the child and do not require a completed fertility. They are also more efficient as they make use of all the information available given the current demographic composition of the family.

We implement our detection tests over a large set of countries, and quantify, for the countries identified by our tests, the magnitude of gender bias in parental preferences. Calibrating a simple model of gender biased preferences, we show that the desired sex ratio exceeds 130 boys for 100 girls in countries such as India, Armenia and Nepal, largely above the actual sex ratios. Overall in the countries in which stopping rule is being practiced, instrumental births represent more than two third of the stopping rule. Studying instrumental births independently of sex-selective abortions can therefore lead to very large underestimation of the prevalence of the stopping rule.

Appendix

2.A Proof of Proposition 1

Let us first assume that the child at rank k is a boy and consider his younger siblings. Three cases arise. In a first case, the desired number of boys is obtained before reaching the maximal number of children, which occur with probability $\sum_{j=1}^{\bar{N}-k-1} B\left(b^{*}-e-1, j\right)$ (where $B\left(a, b\right)$ is the simple binomial probability of having exactly a successes in b trials). In the second case, one needs exactly $\bar{N}$ children to reach the desired number of boys, b^{*}. This occurs with probability $\left(p \sum_{j=1}^{\bar{N}-k-1} B\left(b^{*}-e-2, j\right)\right)$: with their $\bar{N}-1$ younger children, the parents have exactly $n^{*}-1$ boys and, with probability p, their last child, at rank $\bar{N}$, is a boy. Finally, one finds parents who do not reach their desired number of boys when having $\bar{N}$ children.

Consider now a girl of the same rank k who has e older brothers. Suppose first that her next sibling is a boy. For all families that reach their desired number of boys with less than $\bar{N}$ children, this boy will have exactly the same expected number of younger siblings to that of a boy of rank k who has e older brothers. For families which, with a boy at rank k, reach a size $\bar{N}$, his expected number of younger siblings is equal to the expected number of younger siblings of a boy of rank k minus 1. In other words, the expected number of siblings of this boy of rank $k+1$, which we denote by $E\left(Y_b\left(k+1, e\right) \mid g_k\right)$

(to indicate that her sibling of rank k is a girl, g), is given by:

$$
\begin{aligned}
E\left(Y_b\left(k+1,e\right)\mid g_k\right) = {} & E\left(Y_b\left(k,e\right)\right)\left(\sum_{j=1}^{\bar{N}-k-1} B\left(b^{*}-e-1,j\right)\right) \\
& + \left(E\left(Y_b\left(k,e\right)\right)-1\right)\left(1\right. \\
& \left. - \sum_{j=1}^{\bar{N}-k-1} B\left(b^{*}-e-1,j\right)\right) \\
= {} & E\left(Y_b\left(k,e\right)\right)-1 \\
& + \left(\sum_{j=1}^{\bar{N}-k-1} B\left(b^{*}-e-1,j\right)\right)
\end{aligned}
$$

Suppose instead that her next sibling is a girl. Following the same reasoning as above, this girl, of rank $k+1$, has an expected number of younger siblings which is given by:

$$
E\left(Y_g\left(k+1,e\right)\mid g_k\right) = E\left(Y_g\left(k,e\right)\right)-1+\left(\sum_{j=1}^{\bar{N}-k-1} B\left(b^{*}-e,j\right)\right)
$$

As a result, the expected number of younger siblings for a girl of rank k with e older brothers, $E\left(Y_g\left(k,e\right)\right)$,is given by 1 plus expectation of the number of younger siblings of that girl's next sibling:

$$
\begin{aligned}
E\left(Y_g\left(k,e\right)\right) &= 1 + pE\left(Y_b\left(k+1,e\right) \mid g_k\right) \\
&\quad + (1-p)E\left(Y_g\left(k+1,e\right) \mid g_k\right) \\
&= 1 + p\left(E\left(Y_b\left(k,e\right)\right) - 1\right. \\
&\qquad\qquad \left. + \left(\sum_{j=1}^{\bar{N}-k-1} B\left(b^* - e - 1, j\right)\right)\right) \\
&\quad + (1-p)\left(E\left(Y_g\left(k,e\right)\right) - 1\right. \\
&\qquad\qquad \left. + \left(\sum_{j=1}^{\bar{N}-k-1} B\left(b^* - e, j\right)\right)\right) \\
&= E\left(Y_b\left(k,e\right)\right) + \left(\sum_{j=1}^{\bar{N}-k-1} B\left(b^* - e - 1, j\right)\right) \\
&\quad + \frac{1-p}{p}\left(\sum_{j=1}^{\bar{N}-k-1} B\left(b^* - e, j\right)\right)
\end{aligned}
$$

$$
\Longrightarrow E\left(Y_g\left(k,e\right)\right) > E\left(Y_b\left(k,e\right)\right),\ \forall k < \bar{N}, e \leq b^* - 1.
$$

□

2.B The stopping rule with a desired family size

The literature sometimes uses a slightly different approach than the one we follow (see Sheps (1963), for example). While they still assume that parents desire a given number of boys, b^*, parents also have a preference over their total number of children, n^*, which corresponds to their ideal family size. If, with n^* children, they do not have b^* boys, they continue to have children till they reach their desired number of boys. In other words, these parents have lexicographic preferences in n^* and b^*, with $0 < b^* \leq n^*$. To analyze this alternative model, we first assume away a constraint on the maximum number of children so that parents, if needed, have as many children as they need to reach the desired number of boys.

Consider first a family that succeeds in having at least b^* boys with n^*children. In such families, at any rank k, girls and boys have exactly the same number of younger siblings, which is equal to $(n^* - k)$. The proportion of such families in a large population is equal to the probability of having at least b^* 'successes' (boys) in n^* trials (children), which we denote as above $\sum_{j=b^*}^{n^*} B(j, n^*)$. All other families need more than n^*children to reach their desired number of boys. In such families, at any rank k, a girl will have $1/p$ more younger siblings than a boy, $1/p$ corresponding to the expected number of children necessary to have one extra boy. The proportion of such families is given by $\sum_{j=0}^{b^*-1} B(j, n^*)$. We therefore have:

Proposition 3: In families with lexicographic preferences over (b^*, n^*), at any rank, girls have in expected terms

$\left(\frac{1}{p}\sum_{j=0}^{b^*-1} B\left(j,n^*\right)\right)$ more younger siblings than boys of the same rank.

A direct consequence of this proposition is that a girl on average (i.e., over all ranks) will also have $\left(\frac{1}{p}\sum_{j=0}^{b^*-1} B\left(j,n^*\right)\right)$ more younger siblings than a boy. A closer examination of this expression is illustrated in Figure 2.B.1: the difference in the expected number of younger siblings is larger for a smaller desired family size and for a larger desired number of boys. Note for example how, for a given desired number of boys an increase in the ideal family size leads to a decrease in the difference in younger siblings (a change in curve). Note also how, for a given ideal family size, an increase in the number of desired boys increases the difference in younger siblings (a change along the curve). As a result, it is likely that societies undergoing a demographic transition display a stronger differential in younger siblings than societies characterized by larger family sizes, provided the desired number of boys does not vary too much. That is, the fertility squeeze hypothesis (Guilmoto (2009); Jayachandran (2017)) not only applies to sex-selective abortions but also to instrumental births.

Finally, imposing a constraint on family sizes in this setting does not change our main results. Assume again that family size cannot exceed a given level $\bar{N}$. Clearly, this constraint is only binding for families that needed more than n^* children to have their desired number of boys, b^*. Among this subset however, Proposition 1 above applies. More precisely, at any rank $k > n^*$,with $n^* < k < \bar{N}$ and for any number of elder brothers e, with $e \leq b^* - 1$, the expected number of younger siblings is strictly larger for a girl than for a boy.

Figure 2.B.1: Difference in expected number of younger siblings between girls and boys with lexicographic preferences in b^*

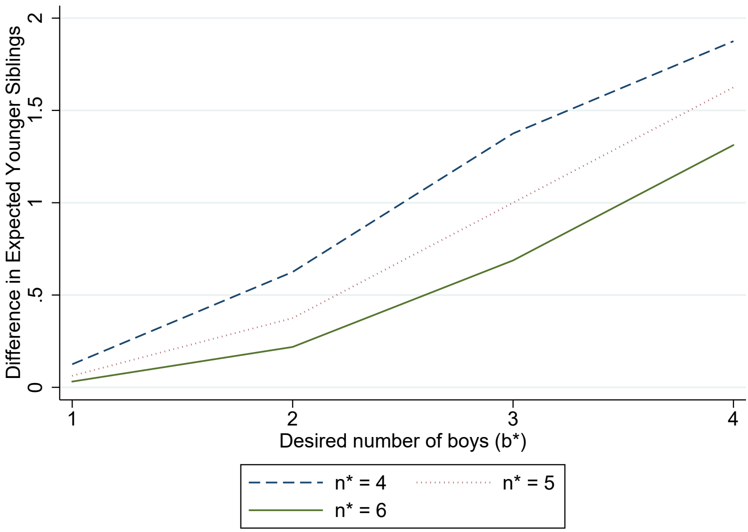

Data Source: Author's simulations.
Reading: When parents have an ideal number of boys b^* of 2, and an ideal family size n^* of 4, girls one average have 0.625 more younger siblings than boys. When parents want the same number of boys but for an ideal family size n^* of 6, girls have on average 0.219 more younger siblings than boys.

2.C Imperfect household information

Because our tests take the perspective of a child and his elder and younger siblings, the ideal dataset to perform our tests is fertility history. However, information on the complete fertility history is not always present in standard household surveys, for instance when the only source of information comes from a household roster. We now discuss the performance of our tests in this setting. A household roster lists the gender, age and family linkage of household members living in the household, excluding children no longer living in this household. A child missing from the household roster has two implications: our measures are not computed for this child (a selection issue) and this child is not accounted for when computing our measures on her siblings (a measurement issue).

As long as the probability of leaving the household is uncorrelated with gender, our tests remain unbiased but are simply less precise. However, in gender biased societies, the presence of a child in a household is correlated with gender, for instance because the age at marriage differs across gender. Consider first the case in which (i) children leave the household upon marriage and (ii) girls marry at a younger age than boys.[45] As a result, relatively older girls are not accounted for when aplying our measures on their siblings. This is also true for older boys, but to a lower extent, given that they leave the household at a later age.

Consider first the detection of instrumental births. Because older boys and girls are unobserved, the average number of younger siblings is biased downwards for both genders (as we apply our measure on younger siblings who, by definition, have less younger siblings than their elder, unobserved, brothers and

[45] Patrilocality, whereby boys do not leave their parents while girls, once married, do, reinforces this bias. In comparison, differential mortality rates among older children are of much lower importance, but our discussion easily extends to this issue as well.

sisters). However, since more elder girls go unobserved (relative to boys), the selection bias is more pronounced for girls and the difference in younger siblings between girls and boys is downwards biased. In terms of measurement bias, because the unobserved children are the elder, their absence does not affect our measure for younger children (i.e. the number of younger siblings of these younger children). The only situation in which our measure is affected is for relatively older boys, who are too young to be themselves married, but whose younger sisters are already married. For these boys, the number of younger siblings we measure is lower than the actual one, which biases upwards the difference in younger siblings between girls and boys. It turns out that, in the numerous simulations we ran, this measurement bias is much less important than the selection bias discussed above. As a result, our test, applied to household rosters, underestimates instrumental births. That is, the bias implied by the use of our test for instrumental births typically do not lead to falsely conclude that they are practiced while they are not (false positive). However, the opposite is true: in presence of such bias, our test can be falsely negative.

We now discuss the detection of sex-selective abortion in this context. As discussed in the main text, under sex-selective abortion, the proportion of girls among older siblings is larger for boys than for girls, and this difference gets larger at later ranks. Since the missing observations in household rosters are older children for which this difference is less important, our measure applied to the observed, later rank, children is upwards biased. In terms of measurement bias, the discussion is more intricate. In general, since more elder girls than boys are missing, the proportion of girls among elder siblings is lower than the actual one. But this underestimation is symmetric across gender and does not, per se, create a bias. The asymmetry is again located in this age interval in which girls tend to be married while boys remain in the household. In the interval, the fact that girls are more often missing does not affect their older brothers, but only their younger brothers or sisters. (For the latter, the proportion of girls among elder siblings we measure is lower and the proportion of boys higher, than the actual ones.) Since boys are more numerous in the interval, and are therefore less often impacted by the disappearance of their younger sisters, the measure is, on average, less biased for boys than for girls. The measurement bias tends therefore to also overestimate the difference in the proportion of girls among elder siblings between boys and girls. In general, household roster surveys overestimate our measure for sex-selective abortion. That is, in presence of such bias, our test for sex-selective abortion may lead to more false positive and less false negative.

Finally, let us consider a survey which only provides the number and the gender composition of children in a household. Absent birth ranks, we cannot reconstruct the number of younger siblings or the gender of older siblings of a particular child. In terms of instrumental births, we can still follow Equation 2.3 and replace the number of younger silbings by the total number of siblings, essentially testing whether girls, on average, live in larger families. While, under gender biased preferences, girls have more younger siblings, which will mechanically translate into a larger number of siblings, girls are also, on average, of lower birth rank, and therefore have fewer older siblings than boys. The latter effect never dominates, and, for large enough sample size, a test based on family size, while less precise, will yield the same outcome as the one proposed in this paper.

2.D List of DHS surveys

Table 2.D lists all the DHS surveys used as well as their number of observations.

	Year of interview	Observations
Afghanistan	2015	125715
Albania	2009	12766
	2017	16128
Angola	2015	42002
Armenia	2000	11286
	2005	10297
	2010	8424
	2016	8771
Azerbaijan	2006	13565
Bangladesh	1994	32590
	1996	29366
	2000	31925
	2007	30527
	2011	45844
	2014	43772
Benin	1996	19359
	2001	19398
	2006	57232
	2012	47152

	2017	45853
Bolivia	1989	22338
	1994	24174
	1998	29473
	2003	45116
	2008	40355
Brazil	1986	12356
	1991	15363
	1996	25513
Burkina Faso	1993	20655
	1999	22145
	2003	41520
	2010	56178
Burundi	1987	11886
	2010	24520
	2016	45419
Cambodia	2000	40990
	2005	40457
	2010	37511
	2014	33290
Cameroon	1998	15187
	2004	29455
	2011	42312
Central Africa	1994	16936
Chad	1997	25739
	2004	21448
	2015	68989
Colombia	1986	11622
	1990	15976
	1995	21830

	2000	21267
	2005	71278
	2010	91399
	2015	62593
Comoros	1996	7913
	2012	11497
Congo	2005	16687
Cote d Ivoire	1994	24870
	1999	7575
	2005	13358
	2012	28211
Dominican Republic	1986	20151
	1991	17168
	1996	19784
	1999	2871
	2007	87585
	2013	77443
Ecuador	1987	11835
Egypt	1988	35519
	1992	38076
	1995	56381
	2000	54780
	2005	61455
	2008	48619
	2014	59266
El Salvador	1985	6383
Ethiopia	1992	44174
	1997	39881
	2003	45540
	2008	41392

Gabon	2000	16878
	2012	23109
Gambia	2013	26601
Ghana	1988	14216
	1993	13298
	1998	13188
	2003	15086
	2008	11888
	2014	23118
Guatemala	1987	14698
	1995	38753
	1999	18581
	2015	55398
Guinea	1999	22943
	2005	27115
	2012	27683
	2018	28887
Guyana	2005	4923
Haiti	1994	12547
	2000	26437
	2006	24830
	2012	29013
	2017	27809
Honduras	2006	50093
India	1993	275172
	1999	268879
	2006	256782
	2015	1315617
	2020	1274250
Indonesia	1987	39719

	1991	74329
	1994	90326
	1997	86276
	2002	79791
	2007	84726
	2012	83650
	2017	86265
Jordan	1990	32812
	1997	24243
	2002	25296
	2007	43460
	2009	38199
	2012	42275
	2017	47040
KE	1989	25173
KE2	1993	23899
KE3	1998	23351
KE4	2003	22074
KE5	2009	22534
KE6	2014	83591
Kazakhstan	1999	8106
Kyrgyzstan	1997	8781
	2012	16180
Lesotho	2004	14708
	2009	14429
	2014	11710
Liberia	1986	17264
	2007	22123
	2013	30804
Madagascar	1992	18931

	1997	21653
	2004	20799
	2009	48464
Malawi	1992	16330
	2000	40421
	2004	35883
	2010	72301
	2015	68074
Maldives	2009	20136
	2017	13922
Mali	1987	12252
	1996	37921
	2001	48407
	2006	52140
	2012	33803
	2018	33379
Mexico	1987	22676
Moldova	2005	9903
Morocco	1987	25518
	1992	22657
	2003	32494
Mozambique	1997	25752
	2003	37443
	2011	37984
Myanmar	2016	22989
Namibia	1992	13372
	2000	14946
	2007	19522
	2013	18090
Nepal	1952	29156

	2001	28955
	2007	26394
	2011	26615
	2017	26028
Nicaragua	2001	34157
Niger	1992	23841
	1998	28888
	2006	34378
	2012	44183
Nigeria	1990	28123
	2003	23038
	2008	104808
	2013	119386
	2018	127545
Pakistan	1991	27369
	2006	39049
	2012	50238
	2018	50495
Paraguay	1990	15346
Peru	1986	13291
	1991	38783
	1996	72390
	2000	65453
	2007	89220
	2012	47261
Philippines	1993	35863
	1998	32626
	2003	30443
	2008	28518
	2013	31680

	2017	47244
Rwanda	1992	19440
	2000	27602
	2005	30072
	2010	32639
	2015	30058
Sao Tome	2008	7620
Senegal	1986	14389
	1993	20815
	1997	27448
	2005	39895
	2011	42510
	2016	22740
Sierra Leone	2008	21136
South Africa	1998	22934
Sri Lanka	1987	17705
Sudan	1990	25805
Swaziland	2006	11410
Tajikistan	2012	19938
	2017	21985
Tanzania	1991	29143
	1996	24890
	1999	11952
	2004	30557
	2010	29777
	2015	37169
Thailand	1987	17803
Timor-Leste	2009	35998
	2016	28682
Togo	1988	10782

	1998	26269
	2014	26264
Trinidad	1987	7837
Tunisia	1988	16463
Turkey	1993	19762
	1998	17791
	2004	22443
	2008	19678
	2013	17871
Uganda	1988	16074
	1995	22752
	2001	23410
	2006	30090
	2011	28609
	2016	57906
Ukraine	2007	8007
Uzbekistan	1996	9650
Vietnam	2002	14383
Yemen	1991	29803
	2013	64602
Zambia	1992	22122
	1996	24799
	2002	23805
	2007	21366
	2013	49207
Zimbabwe	1988	12405
	1994	16777
	1999	14184
	2005	19489
	2010	19279

	2015	20791

2.E "Natural" sex ratios

Table 2.E.1 presents the average sex ratios estimated between 1950 and 1980 by Chao et al. (2019) that we use as reference "natural" sex ratio in our estimation of missing girls at birth.

Table 2.E.1: "Natural" sex ratios from Chao et al. (2019)

	Natural Sex Ratio
Albania	106.37
Armenia	106.26
Azerbaijan	106.24
Bangladesh	105.01
Cameroon	102.71
Colombia	104.73
Comoros	102.97
DR Congo	102.62
Egypt	106.29
Gabon	102.06
India	105.73
Jordan	106.56
Kenya	101.95
Kyrgyzstan	105.27
Nepal	104.82
Niger	104.02
Pakistan	106.24
Rwanda	102.33
Sierra Leone	103.3
Tajikistan	106.22
Turkey	104.69
Yemen	106.16

Chapter 3

Stopping rule and girls mortality: Evidence from South Asian countries.

François Woitrin

Abstract: The *stopping rule* refers to a behaviour by which parents continue child bearing until they reach their desired number of boys. In societies in which gender preferences prevails, this method can heavily affects fertility practices. Under this *stopping rule*, girls have a higher probability of having younger siblings and, therefore, are exposed to more sibling competition. If competition is associated with higher probability to die, this higher competition faced by girls results in higher death rate for girls, even in the absence of any other form of discrimination. We therefore study the level of competition faced by boys and girls, as well as the extent to which competition causes deaths. This allows us to estimates the number of girls who died because of the stopping rule. We find that the *stopping rule* explains up to 20% of girls mortality in the less privileged part of some countries.

3.1 Introduction

In many developing countries, women face a disadvantage compared to men in various aspects of human developement (education, health, etc.). Among the reasons often put forward to explain this pattern is a strong preference for having sons, which results from cultural environment. These cultural factors may include patrilocality (Ebenstein (2014)), old age support (Ebenstein and Leung (2010); Lambert and Rossi (2016)), or the burden of the dowry (Arnold et al. (1998)), among others (see Williamson (1976), Das Gupta et al. (2003), or Jayachandran (2015)). The channel through which these cultural characteristics translate into gender differences in outcome is usually thought to be differential investment from parents into their children (itself caused by differential return on the investments or differential costs of children), meaning parents tend to give less resources to their daughters than to their sons.

In this paper, we build on the intuition of Arnold et al. (1998) and Ray (1998) who suggested that even if they were treated in the same way as their brothers, i.e., even if parents invested as much in their daughters than in their sons, girls may end up worse-off. We explore a specific channel of discrimination, related to the preference for having male offspring, which creates different outcomes for girls without necessarily coming from differential treatment from parents within the household. Indeed, in many countries, the preference for sons manifests itself under the form of a stopping rule: parents continue childbearing until a given number of sons is obtained. This practice has been widely studied and his known to have many consequences in terms of fertility: total fertility may be higher than desired (Sheps (1963)), the birth order of girls within families is, on average, lower for girls than for boys (Basu and de Jong (2010)), and, most importantly, the total number of siblings is higher for girls than for boys (Yamaguchi (1989); Basu and de Jong (2010)). In other words, girls live in larger household. The underlying intuition behind this mechanism is the following: because they continue to have children until they reach their desired number of boys, parents are less likely to continue childbearing after the births of a boy than after the birth of a girl. We argue that these additional siblings, through an increase in competition for resources, may lead to higher mortality rates for girls. Moreover, this additional mortality occurs, independently of the way girls are treated once being born.

We therefore propose a decomposition of the impact of the stopping rule on the gender differential in young-age mortality into two parallel mechanisms. The first mechanism, *indirect discrimination*, characterize the excess mortality caused by the additional siblings of girls if they had been treated in the same way as boys within the household. We call it “indirect” for it originates from fertility preferences (developed before the birth of any specific child) which do not automatically translate into behaviors detrimental to girls. The second mechanism identify the excess mortality caused by the extra competition faced by girls, taking into account that girls may in fact not be treated in the same way as their brothers. Such discrimination may materialize into redistribution of resources to the benefit of boys when competition increases. We call this second mechanism the *direct discrimination* effect of the stopping rule. We add on Jensen (2003) who finds that a significant proportion of the disparities in education between boys and girls can be attributed to fertility behaviors. While he analyses the stopping rule by comparing the education outcomes by gender in different subgroups of households identified on the basis of declared son preference, we use the fertility distortions produced by the stopping rule to infer its impact on mortality at the population level. Unlike Jayachandran and Kuziemko (2011) who study discrimination in breastfeeding and its associated health outcomes in a context of son preference, this paper wants to focus discrimination produced at the population level, while parents do not individually discriminate boys and girls within the household (unlike breastfeeding).

Taking advantage of the DHS surveys for Bangladesh, India, Nepal, and Pakistan, we first discuss and document the more intense competition faced by young girls when the stopping rule prevails. We then estimate, through a survival analysis, the impact of competition on under-five mortality, before measuring the effect of the prevailance of the stopping rule on mortality for each of these countries. We show that this mechanism explains a non-trivial share of mortality among young girls, even in the absence of any form of direct discrimination. This indirect discrimination accounts for close to 2% of total female mortality in India, most of this deaths occurring in the poorest part of the population. This accounts for around 2.5 million deaths in India between 1980 and 2015.

The structure of the paper is as follows: we first present the data, before presenting the theoretical framework and measuring the difference in sibling competition in our four countries in Section 3.3. In Section 3.4 we estimate the effect of sibling competition on mortality which will allow us to compute the excess female mortality generated by the stopping rule.

3.2 The data

We use the Demographic and Health Survey (DHS) from four south Asian countries: Bangladesh, India, Nepal, and Pakistan. These countries were selected for their well-documented strong preferences for boys, while also featuring high mortality rates at young ages[1]. The DHS surveys record the fertility history of ever-married women aged 13 to 49. From these data we know whether and when a child died, which allows us to precisely investigate the link between mortality and fertility outcomes. We select the surveys for which we have information on wealth (this includes the surveys from 2007, 2011, and 2014 in Bangladesh, 2006 and 2025 in India, 2007, 2011, and 2017 in Nepal, and 2006, 2012, and 2018 in Pakistan), and, to ensure comparability across countries, we restrict the sample to children born between 1980 and 2015. We choose this particular time window to maximize the number of observations, while keeping a similar distribution in the years of births in each country. Our main sample consists of 41,883 mothers, and 117,837 children in Bangladesh, of 544,576 mothers, and 1,518,216 children in India, of 24,805 mothers, and 75,816 children in Nepal, and of 32,031 mothers, and 131,093 children in

[1]We chose not to include Afghanistan because the data behaved oddly, making our interpretation unclear

Pakistan.

We also use the official estimates published by the Population Division of the Department of Economic and Social Affairs of the United Nations Secretariat[2]. It presents population estimates from 1950 to the present. From these we retrieve to number of girls born during the period under scrutiny.

3.3 Son preference and fertility

Throughout the developing world, son preference as been shown to be widely prevalent, particularly in South Asia, but also in Eastern Europe, North Africa, or the Near/Middle East. In many of these places, this preference for sons leads parents to continue having children until they reach an ideal number of sons. In the demographic literature, this behavior is known as the ”stopping rule”. The consequences of this practice have been largely studied, particularly in terms of the specific fertility outcomes it generates (e.g. Sheps (1963); Yamaguchi (1989); Clark (2000); Basu and de Jong (2010)).

3.3.1 Demographic consequences of the Stopping Rule

In order to understand the fertility distortions caused by the stopping rule, suppose that parents are only interested in having a certain number b^* of boys and are willing to have as many children as needed to reach that goal. Suppose also

[2] Retrieved 10 April 2023

that, at each birth rank, parents have a given probability p of having a boy[3]. Then, each birth can be seen as a lottery draw in which having a boy is a success, while having a girl is a failure (as she will not contribute to its parents' fertility goal). This "bad" draw therefore needs to be compensated by the birth of a least one additional child, meaning that a girl of a particular rank will have on average exactly the same number of younger siblings as a boy of the same rank, plus the expected number of additional births required to have the boy that she is not. On the other hand, children can not differ in terms of their number of older siblings. Indeed, because there is always a probability p to have a boy and $(1-p)$ to have a girl, at any rank k, both boys and girls have exactly $k-1$ older siblings, and the number of boys among older siblings of children of any gender follows a binomial distribution with $k-1$ trials and probability p of success. As a result, at the population level (i.e. aggregating birth ranks), the only difference in the distribution of siblings of boys and girls comes from the younger siblings.

Table 3.3.1 below illustrates the consequences of the stopping rule for a set of 10,000 households which only want one boy, a maximum family size of four, and with a 50% probability to have a boy at each birth. Suppose also that parents stop their fertility as soon as they get the boy they desired.

[3]Note that the results below hold for all $b^*, g^* \geq 0$ and any finite maximum family size, as long as $b^*/g^* \geq p/(1-p)$. That is, parents desire relatively more boys than their 'natural' prevalence in the population. They also hold for different structure of preferences: parents want a number n^* with at least b^* of boys, parents are ready to have more children to reach b^* than g^*, etc.

In that particular setting, there are five possible household compositions, each represented in a certain proportion in a large population. Although there are as many boys than girls at each birth rank and in the population (the stopping rule can not itself skew the sex ratio, as there is always a probability p to have a boy), girls disproportionately live in the largest households. Indeed, 50% of boys are only-child, while none of the girls are. From the table, one can easily compute the average number of older and younger siblings of both boys and girls. Boys and girls both have an average of 0.733 older siblings[4], boys never have younger siblings while girls have 1.466 on average[5]. As predicted, the reason why girls live in larger household comes from their greater amount of younger siblings. A key assumption of this paper is that the stopping rule is the only mechanism that can distort fertility distributions in such a way. This assumption will allow us to attribute any difference in the number of siblings directly to the stopping rule, without the need to strictly establish causal empirical relationship[6].

[4]Among the 9375 children of each gender, 5000 are first-born and have no older siblings, 2500 have 1, 1250 have 2, and 625 have 3, which makes an average of 0.733 older siblings.

[5]Among the 9375 girls, 625 have no younger siblings (the last-born girls in the fifth category of household), 5000 have only one (2500 girls from the second category of households, 1250 from the third, 625 from the fourth, and 625 from the fifth), 2500 have 2, and 1250 have 3.

[6]Note also that, in the remainder of the paper, we will somewhat imprecisely use to term "stopping rule" to refer to the special case of "gender-biased stopping rule". When parents desire the same number of boys and girls, and continue childbearing until reaching both of these numbers, the stopping rule is gender-unbiased. Total fertility will still be higher than what it would have been if parents were gender-indifferent (parents only

Table 3.3.1: Illustration of the stopping rule

					prop.	#♂	#♀
1.	♂				0.5	5000	0
2.	♀	♂			0.25	2500	2500
3.	♀	♀	♂		0.125	1250	2500
4.	♀	♀	♀	♂	0.0625	625	1875
5.	♀	♀	♀	♀	0.0625	0	2500
Total						9375	9375

Table 3.3.2 illustrates the theory's main prediction with the number of ever-born older and younger siblings at the time of the survey in Indian DHS. There is indeed no difference (at 5% significance level) in the number of older siblings between boys and girls. As expected, the largest part of the observed difference in the number of ever-born siblings comes from the number of younger sibling of girls (1.1617) which is much higher than that of boys (1.0233). Note that the magnitude of these differences depends on the parameters b^*, g^*, p and maximum family size, as well as other factor such as the intensity with which parents practice the stopping rule for instance (are they willing to have as many children in order to reach b^* than to reach g^*). This means that a same difference could be observed for completely different kind of preference profiles. More importantly for our study, it also

desire a specific number of children), but boys and girls will on average have the same number of younger siblings.

means that the difference in the number of siblings that girls have compared to boys will vary from one environment to the other. In particular **?** highlights that, for a given number of desired boys (and girls), the difference in the number of ever-born siblings is monotonically increasing in maximum family size. Conversely, for a given maximum family size, the difference is monotonically decreasing in the number of desired boys.

Table 3.3.2: Stopping Rule and ever-born siblings in India

	Number of Siblings					
	All		Younger		Older	
	♂	♀	♂	♀	♂	♀
Population Average	2.1559	2.2899	1.0233	1.1617	1.1327	1.1282
Difference	.1340***		.1384***		-.0044*	
	(.0032)		(.0024)		(.0025)	

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Finally, it is worth pointing that the practice of sex-selective abortion has developed rapidly, particularly in Asia, as a mean to achieve the desired gender composition in the family (Park and Ho (1995); Arnold et al. (2002); Abrevaya (2009); Jayachandran (2017); Dimri et al. (2019)). In the context of gender-biased preferences, parents can now directly interrupt pregnancies and control the gender of their children. While both mechanisms have the same underlying causes, their demographic consequences are entirely different:

if parents fully control the gender of their children, girls born when sex-selective abortion is possible are the result of a choice (Qian et al. (2014)). Hence, these girls will not necessarily be followed by more pregnancies and will not live in larger families. As a result, one may think that the widespread adoption of sex-selective abortion should eliminate the competition consequences of the stopping rule. While essentially true, it needs to be qualified: when sex-selective abortion does not lead to a perfect control of the gender of the child, but simply affects the relative probability for each gender to be born, young girls will still be exposed to a more intense sibling competition, even though the difference with boys will be mitigated. The major demographic consequence of sex-selective abortion is that it distorts the sex-ratio (which, remember, the stopping rule by itself does not affect on the aggregate). Nevertheless, as we study the effect of the number of siblings and not the effect of their gender composition, the prevalence of sex-selective abortion only affects our mechanism in that it reduces the competition faced by children.

3.3.2 Siblings and Intra-household competition

In order to investigate the effect of those additional siblings through the channel of intra-household competition, we are interested in the number of siblings each child has had to actually compete with, instead of the number of ever-born siblings she had at some point in time. In the remainder of the paper we will talk about 'intra-household competition' or 'sibling competition' to refer to the number of alive siblings any child had to live with during a period of interest. The specific mechanism we have in mind is that if a child has many siblings, his share of a given quantity of resources (food or health care in our context, but education for instance could be another example) will be smaller. We therefore build a measure of siblings competition that is better suited to our analysis. This measure is defined as the average number of siblings, older and younger, per years that a child has been exposed to from his birth, and is constructed as the cumulated number of months during which she was competing with all of her alive siblings. Thus, a sibling born when the child of interest is 18 months old and who died after 10 months, will compete with that child for the 10 months she lived. Moreover, if the child of interest herself dies at 24 months old, they will compete for only 6 months. The cumulated number of months a child was competing with her siblings is then divided by the number of month the child lived. Therefore, if a child who was observed at age 5, had a first younger sibling born when he was 12 months old, and a second when he was 36 months (assuming neither of the two die), she will be in competition with those younger siblings for a total of 72 months, which makes our measure equal to 1.2 : the child has competed with an average of 1.2 siblings per year during his first 5 years. In a similar manner, if a child (also observed at age 5) had two older siblings when he was born, among which one died when he was age 30 months old, our measure of older sibling competition will be equal to 1.5. The fundamental difference between this measure and the theory on ever-born siblings is that mortality is now taken into account so that a child who dies will be considered as only partially competing with his siblings. If the probability for the siblings of girls and boys to die is the same, more ever-born siblings trivially implies more competition[7].

Table 3.3.3 presents the differences in sibling competition (on the living period of each individual at the time of the interview) between gender for each country in our sample. In each country the difference in sibling competition is positive and significant, implying that son preference leads to additional competition for girls. In Nepal for instance, girls face an average of 0.221 additional siblings per year compared to their brothers. Considering the average living-time in the our sample (132 months old), this translates into around 30 cumulated months of additional competition with siblings for girls on average.

Table 3.3.3: Gender differences in Sibling Competition, by country

	Bangladesh	India	Nepal	Pakistan
Diff. in Sibling Competition	0.103***	0.166***	0.221***	0.125***
	(0.011)	(0.003)	(0.014)	(0.014)
Observations	117,826	1,518,216	75,816	123,423

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Wealth will play a crucial role in our analysis for two main reasons. Adding on the obvious impact it has on whether a

[7]If the siblings of girls have a higher mortality rate than the siblings of boys, more ever-born siblings does not imply more competition. This is something that will straightforwardly be taken into account in our analysis

child will actually die from the competition he faces, we highlighted that the differential in the number of ever-born siblings (and therefore in sibling competition) itself depends on fertility preferences. If these preferences are related to wealth, then the differences in sibling competition will also be related to wealth. Household wealth is proxied by a standardized index generated in the DHS with a principal components analysis. The dimensions considered in this analysis may vary from one country to the other, but typically include the ownership of goods (phone, televisions, bicycles, refrigerator for instance) or the living environment of the household (materials used for housing construction, water access or sanitation facilities, among others). This index claims to represents the economic status of the household in a more permanent and easy-to-measure way than income and consumption for instance. In their paper discussing this particular index, **?** find that it correlates with many factors such as mortality or education, and advocate for its use as a proxy for wealth. The continuous-valued index from the PCA is then used to separate households into the five wealth quintiles considered in the remainder of the paper. Figure 3.3.1 presents the differences in competition by wealth in our four countries. In all quintiles of the four countries, girls indeed face more competition from their siblings compared to boys. Breaking down the effect by wealth suggests for instance that the poorest Pakistani girls compete with only 0.067 more siblings than boys on average, while this number is twice as high in other quintiles. Interestingly, in the four countries, the gender differences are smallest in the poorest fraction of the population. While not always large, the differences in the estimates from one quintile to another turns out to be significant, particularly in India and Pakistan.

Figure 3.3.1: Gender difference in Sibling Competition by wealth quintile, by country

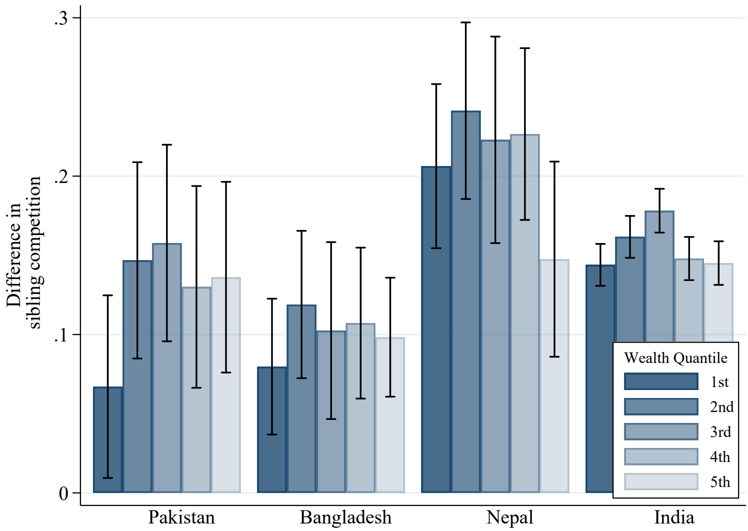

3.4 Sibling competition and mortality

A large literature already documented various forms of discrimination against girls as an important source of differential mortality by gender (Qian (2008); Barcellos et al. (2014)). These "direct" forms of discrimination (investment in childcare, breastfeeding practices, or access to healthcare and nutrition, among others) are not the focus of this paper. Instead, we focus on excess mortality of girls caused by the practice of the stopping rule and its consequences in terms of fertility, which is, at least to some extent, not a direct and deliberate attempt from parents to treat girls and boys differently. In this section, we investigate the impact on mortality of a more intense sibling competition, again defined as the number of alive siblings living with a child of interest.

We therefore run a survival analysis in order to estimate the impact of additional siblings on the probability for a child to die before age 5[8]. Because the number of siblings that a child has typically varies over time, we need to take into account that our variable of interest is not constant. In order to take the time-varying characteristic of our variable into account, we reshape our data so that each child is observed multiple times on four possible occasions: his birth, the birth of (younger) siblings, the death of sibling and his own death (or the date of the interview if not observed). Thanks to the DHS, we know the date (in month) at which all these events occurred. For instance, a child who was born at any time t and who died at age 4, who had a first younger sibling born 10 months after his birth and second 20 months later will be observed 4 times. Instead of using the more popular semi-parametric Cox model with common baseline hazard,

[8]We choose this threshold in order to restrain our analysis on a given period, while including as many deaths as possible. This represents 92.25% of the deaths in Bangladesh, 90.45% in India, 92.31% in Nepal, and 94% in Pakistan)

we specify a parametric survival model[9]. This category of survival models relies on more assumptions in that one needs to specify the underlying hazard function. We choose the Weibull distribution which fits most appropriately survival time in our data allowing us to take into account the variation (decrease in particular) of the hazard rate over time. The parametric model also relaxes the constant hazard ratio hypothesis of the Cox model, which arguably does not fit to our context[10].

Having a child is an endogenous choice and this will create obvious omitted variable issues: individual and parental characteristics may determine both mortality and the number of siblings. To tackle this issue, we propose various child-level variables controlling for mechanisms associating children characteristics to mortality. In particular, we control for birth spacing, which has been shown to have a direct impact on mortality (see Palloni and Millman (1986); Retherford et al. (1989); Jayachandran and Kuziemko (2011); Rossi and Rouanet (2015)), by including dummies indicating whether a sibling was born at least 18 months before, or at least 18 months after the birth of the child of interest. Jayachandran

[9]The main reason is that our within-household analysis (see below) cannot be performed with standard statistical software due to a dimensionality issues. In Stata for instance, Cox models with fixed effects (shared frailty) are fitted using a likelihood function in which the fixed effects are estimated along with the other regression parameters. Each household represents a parameter to estimate, meaning that too many households may exceed the maximum number of estimable parameters. Parametric models, on the other hand, have the benefit that fixed effects are not directly estimated, but instead integrated out of the likelihood.

[10]Babies being typically weaker early on their lives, the impact of competition may not be the same throughout the hazard period.

and Pande (2017) also show that the first born (in particular boys) may be treated differently than the other children. We therefore control for whether a child is the first born of his or her gender. We also include year of birth FE, gender-birth rank FE, as well as a dummy for whether the child is part of a twin birth. Finally, we include household fixed effects which control for all unobserved parental characteristics. A complete description of the variables used in this specification is given in Appendix 3.A. Finally, the standard errors are clustered at the state level, and observations are weighted according to DHS weights.

The remaining difficulty for a credible identification comes from reverse causality. Parents observing a child health condition for instance, may anticipate his death and have more siblings to 'replace' that child. To address the bias this may create, we investigate two potential solutions. The first one is an instrumentation strategy adjusted to survival models as described in**?**. This method is a straightforward two-stage regression approach analogous to the two-stage least squares approach commonly used for IV analysis for linear regression. In order to account for the additional uncertainty from the first-stage estimation, we additionally perform a non-parametric bootstrap to produce more accurate estimates of standard error. For this analysis, we propose the age of the mother at the birth of each child as an instrument for his number of siblings. While correlated with the number of children a mother will still have at the birth of any of her children, this instrument is obviously not exogenous per se. Nevertheless, we argue that conditional on the control variables described above, this instrument meets the exclusion restriction. In particular, the household fixed effect controls for all characteristics of the mother that do not vary throughout her fertility period, so that, within a same household, the age of the mother of each child needs not to directly impact their probability to die. The literature however finds that babies born to adolescent mothers or mother older than 35 (sometimes 40) have a significantly higher risk of death (see **?**, **?**, **?**, **?**, **?**, **?** for instance). Importantly, these studies all focus on mortality up to 6 months and do not discuss similar effects on children on longer periods. The exact threshold for when age directly impact mortality is therefore not a consensus, but it seems reasonable to consider that between 20 and 35, the age of the mother has no consistent significant impact on mortality. Through birth rank FE we take into account mother's experience at taking care of a child, and through birth-spacing we control for potential variations in the time between pregnancies that could occur as the mothers' fertility period shortens.

The second solution we propose is to make use of the panel structure of our data (children being typically observed several times on the period) and of the time-varying nature of our variable of interest to include children fixed effects. We argue that controlling for all time-invariant characteristics of the child (i.e., the known characteristics described above, plus all the unknown ones) captures a consistent part of the remaining endogeneity cause by reverse causality. Note also that this strategy cannot be combined with our instrumentation strategy since our instrument is time-invariant and will therefore be absorbed by the children fixed effects.

Appendix 3.B presents and discuss the results of both specifications in different settings. Because of the instability of the results from the instrumentation strategy arising from the choice of the threshold for the age of the mother at birth, we choose to focus on the results from the specification with children fixed effects. Table 3.4.1 presents these estimated effects for each country in our sample.

Table 3.4.1: Effects of sibling competition on under-five mortality

	Bangladesh	India	Nepal	Pakistan
sibling_competition	0.036***	0.126***	0.033	-0.033**
	(0.013)	(0.006)	(0.020)	(0.015)
Observations	184,207	2,633,671	131,466	255,358

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

These coefficients measure the additional probability to die before age 5 as a consequence of an increase in competition with siblings. For instance, in India, our estimate indicates that one additional sibling in the first 5 years of life increases the probability of death before age 5 by 12.61 percentage points, on average. In Pakistan the effect is negative suggesting than more siblings decreases the probability of dying. Nevertheless, if one considers the possibility that the actual effect of competition depends on other characteristics of the child, looking only at the country-level estimates may not give us the most accurate understanding of what is really happening. In particular, because our mechanism relies on the idea that children with more siblings suffer from more competition for resources, one should expect the effect to depend on the availability of resources themselves, and therefore on the economic status of the household. Figure 3.4.1 displays the results when replicating our first estimations with heterogeneous effects by wealth quintile. As expected, the effect of additional siblings on mortality is consistently decreasing with wealth. Surprisingly, in each country, the effects even get negative in higher quintiles, suggesting that when wealth is high enough additional sibling could be beneficiary to children.

Because the mechanism we study is gender-specific, another important point is whether the effect of siblings on mortality differs by gender. Such difference could have two origins. First, one of the genders could be naturally more sensitive to the scarcity of resources. As infant girls are generally considered as more robust than infant boys, we would expect, if anything, siblings to have a more negative effect on boys. Letting δ_g^* be the "natural" effect of additional competition on girls, and δ_b^* the "natural" effect on boys, we would expect $\delta_g^* \leq \delta_b^*$. The second sources of difference in the effect of competition could be some kind of direct discrimination not captured through our controls and related to competition itself. For instance, when competition increases, parents may reallocate resources between children. In our context of son preference, we expect this mechanism to be at the expense of girls so that $\delta_g > \delta_b$, with δ_g and δ_b the effects of competition on girls' and boys' when taking potential active discrimination mechanisms into account. Given these definitions, we expect $\delta_g > \delta_b^* \geq \delta_g^* > \delta_b$. Running our model with heterogeneous

Figure 3.4.1: Effects of sibling competition on under-five mortality, by wealth quintile

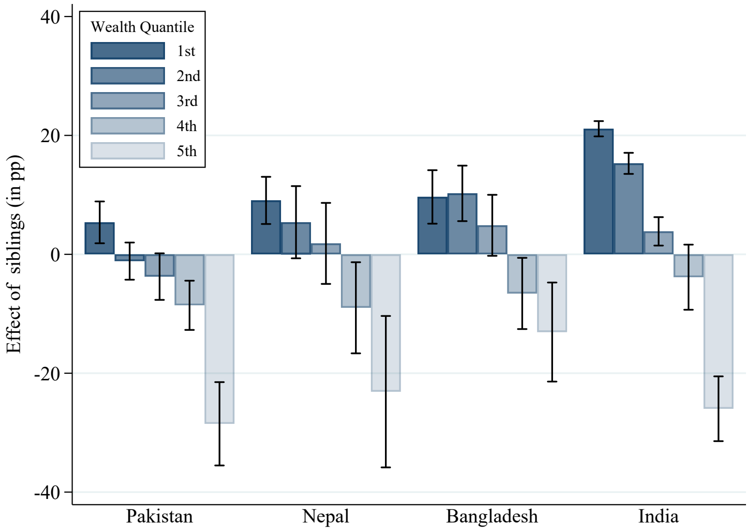

effects by gender would allow us to capture the impact of competition on mortality for each gender separately. A typical way to write the estimated results of a regression is the following:

$$\overline{Mortality_i} = \hat{\delta}_i * \overline{X_i} \qquad \text{with} \qquad i = \{g, b\}$$

where $\overline{Mortality_i}$ is the average mortality for gender i

caused by additional siblings, $\overline{X_i}$ is the average number of siblings in the period of interest, and $\hat{\delta}_i$ is the estimated effects of X in the corresponding group. From the expression above, the BlinderOaxaca method allows us to decompose the gender difference in the mortality resulting from sibling competition:

$$\begin{aligned}\overline{Mortality_g} - \overline{Mortality_b} &= \hat{\delta}_g \overline{X}_g - \hat{\delta}_b \overline{X}_b \\ &= \delta_g^*(\overline{X}_g - \overline{X}_b) + (\hat{\delta}_g - \delta_g^*)(\overline{X}_g - \overline{X}_b) \\ &\quad + (\hat{\delta}_g - \hat{\delta}_b)\overline{X}_b\end{aligned} \tag{3.1}$$

The first term of Equation 3.1 accounts for the excess probability to die from the difference in sibling competition if girls had not been treated differently. This is the measure of "indirect" discrimination we are aiming to estimate: even when parents do not actively discriminate against girls (i.e., they die from sibling competition at the "natural rate" δ_g^*), they still face worse outcomes due to the fertility path followed by their parents. The second term represents the additional probability to die from the difference in sibling competition caused by the stopping rule, taking into account that girls are in fact not necessarily treated in the same way than their brothers. This represents the "direct" discrimination effect of the stopping rule introduced in the Introduction. The last term captures the 'pure' effect of sibling competition if only "active" discrimination was prevalent. That is, even if girls faced the same amount of competition than boys on average, i.e. if the stopping rule was unbiased, they would die more than boys because of differential treatment $\hat{\delta}_g - \hat{\delta}_b$[11]. Because we investigate the mortality consequences of the gender-biased stopping rule, as well as because it represents the combination of numerous effects, each of which deserving to be studied separately, this last term will not be discussed further as part of this paper. It will only be interpreted as the residual effect of sibling competition if the stopping rule had been unbiased. Finally, because we do not have strong arguments suggesting that $\delta_b^* > \delta_g^*$ we simplify our decomposition by considering that $\delta_g^* = \delta_b^* = \delta^*$. Without loss of generality, we suppose hereafter that, in the absence of direct discrimination, the effect of competition is the same for both boys and girls. This hypothesis implies that any significant difference between δ_b^* and δ_g^* can be attributed to active discrimination. Parallelly, it allows us to use the results of Figure 3.4.1, call them $\hat{\delta}^*$, as estimators for the δ^* of each wealth quintile, which could not have been derived otherwise[12]. The differential mortality explained by our passive and active discrimination

[11]Note that this last expression could be decomposed further. Defining $\overline{X}^*$ as the competition boys and girls would have faced without any stopping rule, and $\overline{X}_{usr}$, the competition if the stopping rule had been implemented in an unbiased way (with $\overline{X}_g > \overline{X}_b > \overline{X}_{usr} > \overline{X}^*$), one could have transformed it into an expression which would include terms like $(\overline{X}_b - \overline{X}_{usr})$, or even $(\overline{X}_{usr} - \overline{X}^*)$. This more advanced decomposition would allow us to study the specific impact of implementing the biased stopping rule (compared to its unbiased counterpart), as well as understanding the impact of implementing the stopping rule at all. However, this would only be possible if we knew the counterfactual $\overline{X}_{usr}$ and $\overline{X}^*$, which require to know parents' preferences.

[12]It also makes the results presented below more conservative. One could indeed have argued that when the difference between $\hat{\delta}^*$ and $\hat{\delta}_g$ is not significant, the 'natural' effect of competition for girls is the latter term, taking an upper bound of the results (since $\hat{\delta}_g > \hat{\delta}^*$).

mechanisms are now given by the following expressions:

$$\text{Diff. Mortality}_{\text{passive}} = \delta_g^*(\overline{X}_g - \overline{X}_b) \tag{3.2}$$

$$\text{Diff. Mortality}_{\text{active}} = (\hat{\delta}_g - \delta_g^*)(\overline{X}_g - \overline{X}_b) \tag{3.3}$$

Equation 3.3 highlights that the attribution of deaths to the active discrimination mechanism of the stopping rule relies on a differential effect of sibling competition on mortality between boys and girls. If the difference is not significant, the whole differential mortality that can be attribute to the stopping rule comes from passive discrimination[13]. Note however that an insignificant difference between the mortality effects does not mean that girls are not treated differently compared to boys as $\hat{\delta}_g - \delta_g^*$ is to be interpreted as the differential effect caused by variations in sibling competition. Nevertheless, other sort of differential treatment, linked to other mechanism than sibling competition could still prevail.

Table 3.4.2 presents the results when running the regressions with heterogeneous effects by gender. In India, Bangladesh, and the higher wealth quintiles of Nepal, increasing competition does not have the same impact on boys and

[13] $\hat{\delta}^*$ is a linear combinaison of $\hat{\delta}_g$ and $\hat{\delta}_b$. As a consequence, the classical t-test for the difference between $\hat{\delta}_g$ and $\hat{\delta}^*$ will not be valid as it does not take into account the positive but unknown covariance between the two estimates. Implementing this test will therefore too often fail to conclude in the statistical difference between the effects, making our estimates for the number of deaths from direct discrimination more conservative.

girls. The difference between the effects for both gender is indeed significant, and the difference tends to be increasing (or at least non-decreasing) in wealth, suggesting that wealthier individuals exhibit higher levels of discrimination, other characteristics given. An explanation is that, having more resources to begin with, discrimination can be practice on a larger range, while providing a certain minimum level for each child. This would justify the higher differential, combined with both genders dying less from competition. The consequences of such differential treatment therefore remains of lesser importance as long as girl do not actually die from it.

Table 3.4.2: Effects of sibling competition on under-five mortality, by wealth quintile and gender

Bangladesh

	Child died between 3 and 60 months		
	Girls	Boys	Girls - Boys
Sibling Competition ×			
Wealth Quantile 1	0.120***	0.054	0.070*
	(0.028)	(0.034)	(0.038)
2	0.128***	0.058**	0.073**
	(0.032)	(0.029)	(0.036)
3	0.089***	-0.016	0.108***
	(0.034)	(0.035)	(0.041)
4	-0.018	-0.141***	0.125**
	(0.040)	(0.044)	(0.053)
5	-0.083	-0.208***	0.127**
	(0.051)	(0.060)	(0.063)

India

	Child died between 3 and 60 months		
	Girls	Boys	Girls - Boys
Sibling Competition ×			
Wealth Quantile 1	0.256***	0.159***	0.095***
	(0.008)	(0.009)	(0.010)
2	0.203***	0.095***	0.106***
	(0.010)	(0.012)	(0.014)
3	0.104***	-0.041**	0.142***
	(0.015)	(0.016)	(0.019)
4	0.017	-0.102***	0.116***
	(0.026)	(0.036)	(0.027)
5	-0.181***	-0.348***	0.164***
	(0.034)	(0.038)	(0.047)

Nepal

	Child died between 3 and 60 months		
	Girls	Boys	Girls - Boys
Sibling Competition ×			
Wealth Quantile 1	0.094***	0.086***	0.005
	(0.027)	(0.023)	(0.035)
2	0.066**	0.038	0.023
	(0.033)	(0.053)	(0.072)
3	0.074*	-0.053	0.143**
	(0.042)	(0.045)	(0.060)
4	-0.014	-0.209***	0.223***
	(0.047)	(0.054)	(0.077)
5	-0.166**	-0.313***	0.162*
	(0.074)	(0.080)	(0.096)

Pakistan

	Child died between 3 and 60 months		
	Girls	Boys	Girls - Boys
Sibling Competition ×			
Wealth Quantile 1	0.063***	0.043*	0.035
	(0.019)	(0.023)	(0.027)
2	0.010	-0.035*	0.059**
	(0.019)	(0.021)	(0.028)
3	-0.022	-0.054**	0.047
	(0.027)	(0.023)	(0.036)
4	-0.069***	-0.105***	0.049
	(0.024)	(0.029)	(0.036)
5	-0.335***	-0.244***	-0.070
	(0.048)	(0.042)	(0.058)

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

3.5 Stopping rule and mortality

Using the two decomposition proposed in Equation 3.2 and 3.3, we are now in a position to compute the number of deaths caused by the gender-biased stopping rule. We therefore combine the estimates from Figure 3.4.1 and Table 3.4.2 with a measure of difference in sibling competition in the spirit of the one introduced in Figure 3.3.1. Importantly, our study of mortality being done on a 5-year period, our measure of difference in sibling competition must be computed on the same time-window. However, as the period of interest gets smaller (or finite, in general), our measure becomes increasingly influenced by birth spacing. Indeed, if birth spacing is longer after the birth of a boy than after the birth of a girl, the observed difference at 5 years old will be impacted. As this is a form of discrimination against girls not directly related to fertility behaviors, we look at the difference in sibling competition conditional on birth spacing. We then multiply these expressions with the total number of female births between 1980 and 2015 to estimate the amount of girls who died because of our two mechanisms in each wealth quintile j of each country c:

$$\begin{aligned} \text{Mortality}_{cj} = &\left[\hat{\delta}^*_{cj}(\overline{X}_{gcj} - \overline{X}_{bcj})\right] * \text{Girls}_{cj} \\ &+ \left[(\hat{\delta}_{gcj} - \hat{\delta}^*_{cj})(\overline{X}_{gcj} - \overline{X}_{bcj})\right] * \text{Girls}_{cj} \end{aligned}$$

Table 3.5.1 presents, for each country, the estimated death counts for both mechanisms, with confidence intervals derived from the lower and upper bounds of both constituting effects. The first columns of each category can be interpreted as the minimum number of deaths that each mechanism explains. In India, we estimate that more than 2 million girls died in the poorest 40% of the population as a result of our indirect discrimination mechanism. The estimated number of deaths is smaller for the three other countries, mainly due to the size of their population. In the poorest quintiles of Bangladesh, Nepal and Pakistan, we estimate around 120,000, 25,000 and 90,000 deaths. On the other hand, in the higher wealth quintiles, the number of deaths is negative. This suggests that, in households in which additional siblings have a negative effect on the probability to die, girls are actually the beneficiary of living in larger household. These negative number are therefore represents the number of girls who would have died if they had faced the same amount of competition that their brothers. Note that the numbers of deaths to be attributed to our passive mechanism do not necessarily decline as wealth increase. Indeed, while the effect of mortality was smaller as wealth increased, the difference in sibling competition tended to increase with wealth. This highlights the dual nature of our mechanism: if girls faced less competition, they will not die from indirect discrimination, even if competition increases the probability to die. On the other hand, if girls face intense competition compared to their brothers, but do not die from it, then no excess mortality is to be expected.

Table 3.5.1: Deaths and Stopping Rule

	Number of deaths					
	(min)	**Indirect discrimination**	(max)	(min)	**Direct discrimination**	(max)
Bangladesh						
1	-11,157	**26,512**	108,388	-10,856	**3,643**	58,713
2	612	**51,797**	149,136	-570	**8,381**	85,159
Wealth Quintile 3	-662	**37,526**	127,538	-10,561	**25,392**	135,122
4	762	**-31,486**	-5,730	-128	**31,383**	146,588
5	-33,922	**-77,093**	-48,430	-5,137	**49,080**	203,029
India						
1	869,786	**1,131,530**	1,418,303	-9,675	**86,171**	217,570
2	832,792	**1,082,099**	1,363,888	78,156	**260,516**	487,109
Wealth Quintile 3	97,134	**290,157**	526,715	272,441	**561,240**	910,925
4	-498,923	**-240,549**	116,495	49,828	**498,729**	1,074,585
5	-1,455,835	**-1,417,121**	-1,288,737	260,660	**731,562**	1,330,542
Nepal						
1	733	**12,295**	33,473	-983	**-636**	15,143
2	-551	**10,837**	36,638	-6,061	**2,186**	30,557
Wealth Quintile 3	-4,446	**3,783**	27,925	-3,258	**13,783**	54,961
4	-25,472	**-23,813**	-5,017	-3,428	**25,318**	80,467
5	-15,069	**-33,822**	-26,000	-3,831	**13,851**	70,291
Pakistan						
1	14,189	**89,147**	227,066	-15,121	**50,434**	205,938
2	-34,323	**-19,420**	51,271	-4,193	**73,736**	238,572
Wealth Quintile 3	-55,822	**-63,461**	4,357	-19,223	**67,116**	280,796
4	-43,009	**-113,605**	-102,594	-7,996	**50,428**	230,573
5	-65,964	**-297,401**	-408,596	-22,498	**-7,279**	203,735

Figure 3.5.1 presents the mortality rates by country, i.e., the probability at birth to die from our mechanisms. In India for instance, around 1.7% of the girls in the 2 poorest quintiles die because of the indirect effect of the stopping rule. This number goes down in wealthier household. In the poorest quintiles of the other countries, this probability is constant around 0.5%. While there is no direct discrimination effect on mortality in Bangladesh, Nepal, and Pakistan, India interestingly displays a switch in the relative prevalence of each mechanism as wealth increases. This is because the effect of sibling competition strongly decreases as wealth increases, while, as briefly discussed before, the resource reallocation effect tends to increase with wealth.

Figure 3.5.1: Probability of dying because of the stopping rule.

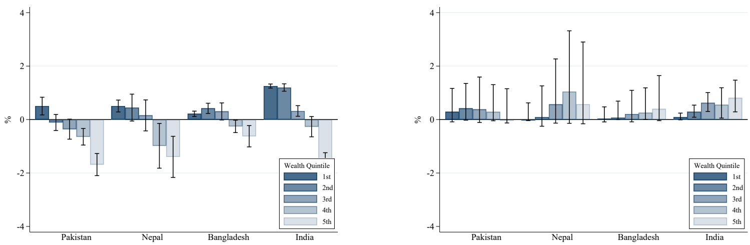

(a) Indirect discrimination mechanism

(b) Direct discrimination mechanism

Figure 3.5.2 depicts the share of deaths that our mechanisms explain. That is the number of deaths relative to the total number of girls who were born between 1980 and 2015 and who died (at any point in time) during this period. In India, the passive discrimination mechanism explains almost 20% of deaths in the lowest wealth quintiles, while reaching half of that number in the other countries. The contribution to mortality of the stopping rule through sibling competition is therefore substantial.

Figure 3.5.2: Share of deaths caused by the stopping rule.

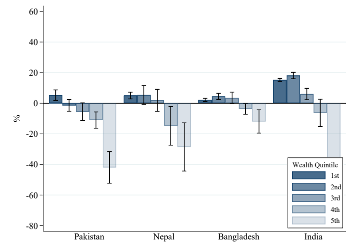

(a) Indirect discrimination mechanism

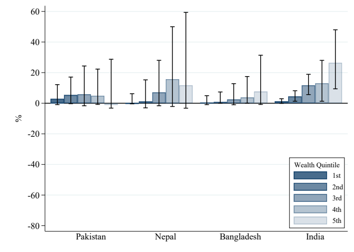

(b) Direct discrimination mechanism

3.6 Conclusion

In many countries, young girls face higher mortality rates than boys. This disadvantage partly follows from active discrimination in the family in terms of access to health care or essential resources. In this paper, we explore a specific channel that explains part of the higher mortality of girls, some in the absence of active and direct discrimination, showing that indirect or implicit discrimination also plays an important role. In an environment where male children are strongly preferred, parents tend to keep having children until they reach their desired number of boys, a behavior known as the stopping rule. As a result, girls have on average more younger siblings than their brothers and therefore experience a stronger intra-household competition for resources. This means that even if girls were not discriminated against after their births, they still die more because of the higher competition they face. Between 1980 and 2015, we estimate over 2,500,000 deaths of girls directly linked to our mechanism in India, 120,000 in Bangladesh, 90,000 in Pakistan, and 25,000 in Nepal. These numbers account for a significant proportion of girls mortality in South Asia, the vast majority of which in the poorest part of the population.

Appendix

3.A Controls variables and descriptive statistics

India

	Mean	S.d.	Min	Max
Child died between 3 and 60mo	0.04	0.20	0	1
Younger siblings competition	0.43	0.43	0	5
Age of the mother at birth	22.90	4.99	0	49
Older siblings competition	1.15	1.34	0	15
Time to next child < 18 months	0.13	0.33	0	1
Time from previous child < 18 months	0.12	0.33	0	1
Girls	0.48	0.50	0	1
Birth rank	2.46	1.60	1	17
Child is from twin birth	0.01	0.11	0	1
Observations	1694986			

Bangladesh

	Mean	S.d.	Min	Max
Child died between 3 and 60mo	0.07	0.25	0	1
Younger siblings competition	0.34	0.38	0	3
Age of the mother at birth	21.81	5.54	10	46
Older siblings competition	1.21	1.46	0	13
Time to next child < 18 months	0.10	0.29	0	1
Time from previous child < 18 months	0.09	0.29	0	1
Girls	0.49	0.50	0	1
Birth rank	2.73	1.86	1	20
Child is from twin birth	0.01	0.12	0	1
Observations	172064			

Nepal

	Mean	S.d.	Min	Max
Child died between 3 and 60mo	0.08	0.27	0	1
Younger siblings competition	0.40	0.39	0	2
Age of the mother at birth	23.48	5.28	9	45
Older siblings competition	1.21	1.42	0	11
Time to next child < 18 months	0.11	0.31	0	1
Time from previous child < 18 months	0.11	0.31	0	1
Girls	0.49	0.50	0	1
Birth rank	2.76	1.82	1	16
Child is from twin birth	0.01	0.12	0	1
Observations	107431			

Pakistan

	Mean	S.d.	Min	Max
Child died between 3 and 60mo	0.05	0.21	0	1
Younger siblings competition	0.62	0.52	0	4
Age of the mother at birth	24.64	5.64	9	46
Older siblings competition	1.80	1.90	0	13
Time to next child < 18 months	0.22	0.42	0	1
Time from previous child < 18 months	0.21	0.41	0	1
Girls	0.48	0.50	0	1
Birth rank	3.28	2.18	1	19
Child is from twin birth	0.02	0.13	0	1
Observations	127983			

3.B Different specifications of the main regression

The first four tables displays the results for different specifications of the model presented in Section 3.4 for each country. The first columns present the results of the survival model with the ”basic” controls, while the second columns show the results of the specification with children fixed effects instead of the controls at the child level. The third columns present the results for the IV regression on the full samples, and the last columns shows the results for the IV regression on the sub-sample of children whose mother was between 20 and 35 years old at the time of their births.

Table 3.B.1: Specifications

	Bangladesh			
	Mortality			
	Survival	Survival	Survival IV	Survival IV (20-35)
# of siblings	-0.012	0.036***	-0.020	0.134***
	(0.028)	(0.013)	(0.115)	(0.130)
First-Stage F-Stat			960	735
Observations	184207	184207	184207	106665
Child FE	No	Yes	No	No
Mother FE	Yes	No	Yes	Yes
Year of Birth FE	Yes	No	Yes	Yes
First of Gender FE * Gender FE	Yes	No	Yes	Yes
Birth rank FE * Gender FE	Yes	No	Yes	Yes
Birth Spacing previous	Yes	No	Yes	Yes
Birth Spacing next	Yes	No	Yes	Yes
Twin FE	Yes	No	Yes	Yes

	India			
	Mortality			
	Survival	Survival	Survival IV	Survival IV (20-35)
# of siblings	-0.111***	0.126***	-0.076***	0.011*
	(0.010)	(0.006)	(0.028)	(0.035)
First-Stage F-Stat			10,454	8,204
Observations	2633671	2633671	2633671	1911299
Child FE	No	Yes	No	No
Mother FE	Yes	No	Yes	Yes
Year of Birth FE	Yes	No	Yes	Yes
First of Gender FE * Gender FE	Yes	No	Yes	Yes
Birth rank FE * Gender FE	Yes	No	Yes	Yes
Birth Spacing previous	Yes	No	Yes	Yes
Birth Spacing next	Yes	No	Yes	Yes
Twin FE	Yes	No	Yes	Yes

	Nepal			
	Mortality			
	Survival	Survival	Survival IV	Survival IV (20-35)
# of siblings	-0.120***	0.033	-0.052***	-0.001
	(0.030)	(0.020)	(0.099)	(0.113)
First-Stage F-Stat			707	567
Observations	131466	131466	131466	93897
Child FE	No	Yes	No	No
Mother FE	Yes	No	Yes	Yes
Year of Birth FE	Yes	No	Yes	Yes
First of Gender FE * Gender FE	Yes	No	Yes	Yes
Birth rank FE * Gender FE	Yes	No	Yes	Yes
Birth Spacing previous	Yes	No	Yes	Yes
Birth Spacing next	Yes	No	Yes	Yes
Twin FE	Yes	No	Yes	Yes

	Pakistan			
	Mortality			
	Survival	Survival	Survival IV	Survival IV (20-35)
# of siblings	-0.299***	-0.033**	-0.077***	-0.0109
	(0.022)	(0.015)	(0.070)	(0.083)
First-Stage F-Stat			640	761
Observations	255358	255358	255358	198266
Child FE	No	Yes	No	No
Mother FE	Yes	No	Yes	Yes
Year of Birth FE	Yes	No	Yes	Yes
First of Gender FE * Gender FE	Yes	No	Yes	Yes
Birth rank FE * Gender FE	Yes	No	Yes	Yes
Birth Spacing previous	Yes	No	Yes	Yes
Birth Spacing next	Yes	No	Yes	Yes
Twin FE	Yes	No	Yes	Yes

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The two specifications which were potential candidates (columns 2 and 4) always agree on the sign of the effect but the differences between them are quantitatively important, although never statistically significant (standard errors from the IV being quite big). The decision of choosing the basic survival model with children fixed effects instead of the IV was made because of the huge variability of the IV results over the choice of the ages at which the sample was censored . Table 3.B.2 below replicates the results of the IV for several plausible minimum (rows) and maximum (columns) values of age at which the age of your mother at birth has no impact on your mortality. This variation is either because our instrument is not strong enough. While significant, most variation in the First-Stage is captured by the household and the birth rank fixed effects with which the age of the mother is obviously strongly correlated. Its remaining explanatory power is therefore relatively low. Moreover, being time-invariant, this instrument is unable to capture the variation in time of the number of siblings. This could also simply be that the instrument, despite our controls and censoring, is not indeed exogenous. Note that for each country, the estimate from the model with children fixed effect is inside the range of the IV results, which encourage the use of this specification. For instance, inn Bangladesh, it closely matches the IV results when selecting mothers between 20 and 33. In India, Nepal, and Pakistan, the estimates from the children fixed effects regressions matches the IV estimates on the selection of mothers between 22 and 33, 21 and 36, and 20 and 34 respectively.

Table 3.B.2: Results from IV with censored samples

		Estimate				
		33	34	35	36	37
Bangladesh						
	18	0.0145	0.0666	0.1090	0.1555	0.1459
	19	0.0616	0.0492	0.0946	0.1267	0.1360
	20	0.0367	0.0845	0.1344	0.1686	0.1845
	21	-0.0180	0.0604	0.1217	0.1625	0.1735
	22	0.0524	0.1125	0.1808	0.2235	0.2371
India						
	18	-0.0412	-0.0199	-0.0558	0.0132	0.0290
	19	-0.0129	0.0096	-0.0229	0.0442	0.0113
	20	0.0213	0.0453	0.0106	0.0365	0.0999
	21	0.0524	0.0291	0.0463	0.0641	0.0856
	22	0.1171	0.1435	0.1009	0.1169	0.1411
Nepal						
	18	-0.1128	-0.1001	-0.0542	-0.0443	-0.0487
	19	-0.0858	-0.0734	-0.0232	-0.0129	-0.0190
	20	-0.0691	-0.0573	-0.0017	0.0092	0.0014
	21	-0.0661	-0.0525	0.0108	0.0232	0.0135
	22	-0.0174	-0.0041	0.0665	0.0796	0.0656
Pakistan						
	18	-0.1888	-0.1431	-0.1290	-0.1097	-0.0725
	19	-0.1534	-0.1139	-0.0853	-0.0261	-0.0350
	20	-0.0945	-0.0394	-0.0190	0.0426	0.0338
	21	-0.0907	-0.0456	-0.0123	0.0104	0.0436
	22	-0.0788	-0.0237	0.0133	0.0393	0.0594

Bibliography

Abdulkadiroğlu, A, YK Che, and Yosuke Yasuda, "Resolving conflicting preferences in school choice : The " Boston mechanism " reconsidered," *American Economic Review*, 2011, *101* (1), 399–410.

Abdulkadiroglu, Atila, Joshua Angrist, Yusuke Narita, and Parag A. Pathak, "Breaking Ties: Regression Discontinuity Design Meets Market Design," *Discution paper*, 2019, (2170).

Abrevaya, Jason, "Are There Missing Girls in the United States? Evidence from Birth Data," *American Economic Journal: Applied Economics*, 2009, *1* (2).

Almond, Douglas and Lena Edlund, "Son-biased sex ratios in the 2000 United States Census," *Proceedings of the National Academy of Sciences*, 2008, *105* (15), 5681–5682.

Altindag, Onur, "Son Preference, Fertility Decline, and the Nonmissing Girls of Turkey," *Demography*, 2016, *53*, 541–566.

Anderson, S. and D. Ray, "Missing Women: Age and Disease," *Review of Economic Studies*, 2010, *77*, 1262–1300.

Andersson, Tommy, Lars Ehlers, and Lars-Gunnar Svensson, "Least manipulable envy-free rules in economies with indivisibilities," *Mathematical Social Sciences*, May 2014, *69*, 43–49.

Anukriti, S, Sonia Bhalotra, and Eddy H.F. Tam, "On the Quantity and Quality of Girls: Fertility, Parental Investments and Mortality," *Economic Journal*, January 2022, *132* (641), 1–36.

Arnold, F., M. K. Choe, and T. K. Roy, "Son Preference, the Family-Building Process and Child Mortality in India," *Population Studies*, 1998, *52(3)*, 301315.

_ , S. Kishor, and T. K. Roy, "Sex-Selective Abortions in India," *Population and Development Review*, 2002, *28(4)*, 759785.

Arnold, Fred, "Measuring the Effect of Sex Preference on Fertility: the Case of Korea," *Demography*, 1985, *22* (2).

_ , "Gender Preferences for Children," *Demographic and Health Surveys Comparative Studies*, 1997, *23.*

Arribillaga, R. Pablo and Jordi Massó, "Comparing Generalized Median Voter Schemes According to Their Manipulability," *Theoretical Economics*, 2015, *11* (2), 547–586.

Ashraf, Nava, Nathalie Bau, Nathan Nunn, and Alessandra Voena, "Brideprice and Female Education," *Journal of Political Economy*, 2020, *128* (2).

B., Moretti Enrico Dahl Gordon, "The Demand for Sons," *Review of Economic Studies*, 2008, *75*, 1085–1120.

Barcellos, Silvia Helena, Leandro S. Carvalho, and Adriana Llenas-Muney, "Child Gender and Parental Investments in India: Are Boys and Girls Treated Differently?," *American Economic Journal: Applied Economics*, 2014, *6* (1).

Basu, D. and R. de Jong, "Son Targeting Fertility Behavior: Some Consequences and Determinants," *Demography*, 2010, *47(2)*, 521536.

Ben-Porath, Yoram and Finis Welch, "Do Sex Preferences Really Matter?," *Quarterly Journal of Economics*, 1976, *90* (2).

Bhalotra, Sonia and Arthur van Soest, "Birth-spacing, fertility and neonatal mortality in India: Dynamics, frailty, and fecundity," *Journal of Econometrics*, 2008, *143* (2), 274–290.

_ **and T Cochrane**, "Where Have All the Young Girls Gone? Identifying Sex-Selective Abortion in India.," *IZA Discussion Paper 5381*, 2010.

Bhaskar, V., "Sex Selection and Gender Balance," *American Economic Journal: Microeconomics*, February 2011, *3* (1), 214–44.

Bongaarts, John, "The Implementation of Preferences for Male Offrspring," *Population and Development Review*, 2013, *39* (2).

Bonkoungou, Somouaoga and Alexander Nesterov, "Reforms meet fairness concerns in school and college admissions," *Working paper*, 2020.

Bruckner, Tim A, Ralph Catalano, and Jennifer Ahern, "Male fetal loss in the U.S. following the terrorist attacks of September 11, 2001," *BMC Public Health*, 2010, *10* (273).

Calsamiglia, Caterina, Guillaume Haeringer, and Flip Klijn, "Constrained school choice: an experimental study," *American Economic Review*, 2010, *100* (4), 1860–74.

Cassan, Guilhem and Lore Vandewalle, "Identities and public policies: Unexpected effects of political reservations for women in India," *World Development*, 2021, *143*, 105408.

Catalano, Ralph and Tim Bruckner, "Economic antecedents of the Swedish sex ratio," *Social Science and Medicine*, 2005, *60* (3), 537–543.

_ , _ , and Kirk R. Smith, "Ambient temperature predicts sex ratios and male longevity," *Proceedings of the National Academy of Sciences*, 2008, *105* (6), 2244–2247.

Chahnazarian, Anouch, "Determinants of the sex ratio at birth: Review of recent literature," *Biodemography and Social Biology*, 1988, *35* (3-4), 214–235.

Chakravarti, Uma, "Conceptualising Brahmanical Patriarchy in Early India: Gender, Caste, Class and State," *Economic and Political Weekly*, 1993, *28* (14), 579–585.

Chao, Fengqing, Patrick Gerland, Alex R. Cook, and Leontine Alkema, "Systematic assessment of the sex ratio at birth for all countries and estimation of national imbalances and regional reference levels," *Proceedings of the National Academy of Sciences*, 2019, *116* (19), 9303–9311.

Chen, Yan and Onur Kesten, "Chinese college admissions and school choice reforms: A theoretical analysis," *Journal of Political Economy*, 2017, *125* (1), 000–000.

Chen, Yuyu, Hongbin Li, and Lingsheng Meng, "Prenatal Sex Selection and Missing Girls in China: Evidence from the Diffusion of Diagnostic Ultrasound," *The Journal of Human Resources*, 2013, *48* (1), 36–70.

Clark, S., "Son Preference and Sex Composition of Children: Evidence from India," *Demography*, 2000, *37(1)*, 95–108.

Combe, Julien, Olivier Tercieux, and Camille Terrier, "The Design of Teacher Assignment: Theory and Evidence," *Working paper*, 2017.

Corno, Lucia, Nicole Hildebrandt, and Alessandra Voena, "Age of Marriage, Weather Shocks and the Direction of Marriage Payments," *Econometrica*, 2020, *88* (3).

Darnovsky, Marcy, "Countries with laws or policies on sex selection," *memo prepared for the April 13 New York City sex selection meeting*, 2009.

Dasgupta, Partha and Eric Maskin, "On the robustness of majority rule," *Journal of the European Economic Association*, September 2008, *6* (5), 949–973.

Decerf, Benoit and Martin Van der Linden, "In search of advice for participants in constrained school choice," *SSRN Working Paper*, 2018, *No. 3100311.*

_ **and Martin Van der Linden**, "Manipulability in constrained school choice," *Available at SSRN 2809566*, 2018.

Dimri, Aditi, Veronique Gille, and Philip Ketz, "Measuring sex-selective abortion: Are there repeated abortions?," *Working Paper*, 2019.

Dogan, Battal and Lars Ehlers, "Minimally Unstable Pareto Improvements over Deferred Acceptance," *SSRN working paper*, 2020.

_ and _ , "Robust Minimal Instability of the Top Trading Cycles Mechanism," *SSRN working paper*, 2020.

Dubuc, Sylvie and Devinderjit Singh Sivia, "Is sex ratio at birth an appropriate measure of prenatal sex selection? Findings of a theoretical model and its application to India," *BMJ Global Health*, 2018, *3.*

Ebenstein, Avraham, "Patrilocality and Missing Women," *Working Paper*, 2014.

_ and Steven Leung, "Son Preference and Access to Social Insurance: Evidence from China's Rural Pension Program," *Population and Development Review*, 2010, *36* (1), 47–70.

Edlund, Lena, Hongbin Li, Junjian Yi, and Junsen Zhang, "Sex Ratios and Crime: Evidence from China," *The Review of Economics and Statistics*, 12 2013, *95* (5), 1520–1534.

Ergin, Haluk and Tayfun Sönmez, "Games of school choice under the Boston mechanism," *Journal of Public Economics*, January 2006, *90* (1-2), 215–237.

Field, Erica, Seema Jayachandran, and Rohini Pande, "Do Traditional Institutions Constrain Female Entrepreneurship? A Field Experiment on Business Training in India," *American Economic Review Papers and Proceedings*, 2010, *100* (2), 125–129.

Filmer, Deon, Jed Friedman, and Norbert Schady, "Development, Modernization, and Childbearing: The Role of Family Sex Composition," *World Bank Economic Review*, 2009, *23* (3), 371–398.

Fleurbaey, Marc, "Social preferences for the evaluation of procedures," *Social Choice and Welfare*, 2012, *39*, 599–614.

Gerber, Anke and Salvador Barberà, "Sequential voting and agenda manipulation," *Theoretical Economics*, 2016, *Forthcoming.*

Goodkind, Daniel, "On Substituting Sex Preference Strategies in East Asia: Does Prenatal Sex Selection Reduce Postnatal Discrimination?," *Population and Development Review*, 1996, *22* (1), 111–125.

Grosjean, Pauline and Rose Khattar, "Its Raining Men! Hallelujah? The Long-Run Consequences of Male-Biased Sex Ratios," *The Review of Economic Studies*, 05 2018, *86* (2), 723–754.

Guilmoto, Christophe, "The sex ratio transition in Asia," *Population and Development Review*, 2009, *35* (3).

Gupta, Monica Das, "Is banning sex-selection the best approach for reducing prenatal discrimination?," *Asian Population Studies*, 2019, *15* (3), 319–336.

— , **Jiang Zhenghua, Li Bohua, Xie Zhenming, Woojing Chung, and Bae Hwa-Ok**, "Why is Son Preference so Persistent in East and South Asia? A Cross-Country Study of China, India and the Republic of Korea," *Journal of Development Studies*, 2003, *40* (2), 153–187.

Haeringer, Guillaume and Flip Klijn, "Constrained school choice," *Journal of Economic Theory*, September 2009, *144* (5), 1921–1947.

Haughton, Jonathan and Dominique Haughton, "Are Simple Tests of Son Preference Useful? An Evaluation Using Data from Vietnam," *Journal of Population Economics*, 1998, *11* (4), 495–516.

Helle, Samuli, Samuli Helama, and Kalle Lertola, "Evolutionary ecology of human birth sex ratio under the compound influence of climate change, famine, economic crises and wars," *Journal of Animal Ecology*, 2009, *78* (6), 1226–1233.

Hesketh, Therese and Zhu Wei Xing, "Abnormal sex ratios in human populations: Causes and consequences," *PNAS*, 2006, *103* (36).

Hu, Luojia and Analía Schlosser, "Prenatal Sex Selection and Girls WellBeing: Evidence from India," *The Economic Journal*, 08 2015, *125* (587), 1227–1261.

Jayachandran, S., "The Roots of Gender Inequality in Developing Countries," *Annual Review of Economics*, 2015, *7*, 63–88.

— , "Fertility Decline and Missing Women," *American Economic Journal: Applied Economics*, 2017, *9* (1).

— **and I. Kuziemko**, "Why Do Mothers Breastfeed Girls Less than Boys? Evidence and Implications for Child Health in India," *The Quarterly Journal of Economics*, 2011, *126(3)*, 1485–1538.

— **and R. Pande**, "Why Are Indian Children So Short? The Role of Birth Order and Son Preference," *American Economic Review*, 2017, *107* (9), 2600–2629.

Jensen, R., "Equal Treatment, Unequal Outcomes? Generating Sex Inequality through Fertility Behaviour," 2003. Working Paper, Mimeo, Harvard University.

Jha, P., M. A. Kesler, R. Kumar, F. Ram, U. Ram, L. Aleksandrowicz, D. G. Bassani, S. Chandra, and J. K. Banthia, "Trends in selective abortions of girls in India: analysis of nationally representative birth histories from 1990 to 2005 and census data from 1991 to 2011," *Lancet*, 2011, *377*, 1921–1928.

Jha, Prabhat, Rajesh Kumar, Priya Vasa, Neeraj Dhingra, Deva Thiruchelvam, and Rahim Moineddin, "Low male-to-female sex ratio of children born in India: national survey of 1ů1 million households," *The Lancet*, 2006, *367* (9506), 211–218.

Kalsi, Priti, "Abortion Legalization, Sex Selection, and Female University Enrollment in Taiwan," *Economic Development and Cultural Change*, 2015, *64* (1), 163–185.

Kapadia, Karin, "Mediating the Meaning of Market Opportunities: Gender, Caste and Class in Rural South India," *Economic and Political Weekly*, 1997, *32* (52), 3329–3335.

Klijn, Flip, Joana Pais, and Marc Vorsatz, "Preference intensities and risk aversion in school choice: a laboratory experiment," *Experimental Economics*, 2013, *16*, 1–22.

Lambert, Sylvie and Pauline Rossi, "Sons as widowhood insurance: Evidence from Senegal," *Journal of Development Economics*, 2016, *120*, 113–127.

Lin, Min-Jeng, Jin-Tan Liu, and Nancy Qian, "More Missing Women, Fewer Dying Girls: The Impact of Abortion on Sex Ratios at Birth and Excess Female Mortality in Taiwan," *Journal of the European Economic Association*, 2014, *12* (4).

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

McClelland, Gary H., "Determining the Impact of Sex Preferences on Fertility: A Consideration of Parity Progression Ratio, Dominance, and Stopping Rule Measures," *Demography*, 1979, *16* (3).

Milazzo, Annamaria, "Why are adult women missing? Son preference and maternal survival in India," *Journal of Development Economics*, September 2018, *134*, 467–484.

Mohapatra, Seema, "Global Legal Responses to Prenatal Gender Identification and Sex Selection," 2013, *690.*

Nandi, Arindam and Anil B. Deolalikar, "Does a legal ban on sex-selective abortions improve child sex ratios? Evidence from a policy change in India," *Journal of Development Economics*, 2013, *103*, 216–228.

Norling, Johannes, "A New Framework for Measuring Heterogeneity in Childbearing Strategies When Parents Want Sons and Daughters," *Working Paper*, 2015.

Palloni, Alberto and Sara Millman, "Effects of Inter-Birth Intervals and Breastfeeding on Infant and Early Childhood Mortality," *Population Studies*, 1986, *40*, 215–236.

Park, Chai Bin, "The Fourth Korean Child: The Effect of Son Preference on Subsequent Fertility," *Journal of Biosocial Science*, 1978, *10*, 95–106.

_ , "Preference for Sons, Family Size, and Sex Ratio: An Empirical Study in Korea," *Demography*, August 1983, *20* (3).

_ **and Nam-Hoon Cho**, "Consequences of Son Preference in a Low-Fertility Society: Imbalance of the Sex Ratio at Birth in Korea," *Population and Development Review*, 1995, *21* (1), 59–84.

_ **and Nam-Hoon Ho**, "Consequences of Son Preference in a Low-Fertility Society: Imbalance of the Sex Ratio at Birth in Korea," *Population and Development Review*, 1995, *21* (1), 59–84.

Pathak, Parag A and Tayfun Sönmez, "School admissions reform in chicago and england : comparing mechanisms by their vulnerability to manipulation," *American Economic Review*, 2013, *103* (1), 80–106.

Qian, Nancy, "Missing Women and the Price of Tea in China: The Effect of Sex-Specific Income on Sex Imbalance," *Quarterly Journal of Economics*, 2008, *123* (3), 1251–1285.

_ , **Ming-Jen Lin, and Jin-Tan Liu**, "More Missing Women, Fewer Dying Girls: The Impact of Abortion on Sex Ratios at Birth and Excess Female Mortality in Taiwan," *Journal of the European Economic Association*, 2014, *12* (4), 899–926.

Ray, Debraj, *Development Economics*, Princeton University Press, 1998.

Retherford, Roert D., Minja Kim Choa, Shyam Thapa, and Bhakta B. Gubhaju, "To What Extent Does Breast-feeding Explain Birth-Interval Effects on Early Childhood Mortality?," *Demography*, 1989, *26* (3), 439–450.

Rosenblum, Daniel, "The Effect of Fertility Decisions on Excess Female Mortality in India," *Journal of Population Economics*, 2013, *26(1)*, 147–180.

Rossi, Pauline and Lea Rouanet, "Gender Preferences in Africa: A Comparative Analysis of Fertility Choices," *World Development*, 2015, *72*, 326–345.

Selten, Reinhard, "Properties of a measure of predictive success," *Mathematical social sciences*, 1991, *21* (2), 153–167.

Sen, A., "More than 100 Million Women are Missing," *New York Review of Books*, 1990, *37(20)*, 61–66.

Sheps, Mindel C., "Effects on Family Size and Sex Ratio of Preferences Regarding the Sex of Children," *Population Studies*, 1963, *17* (1).

Tuljapurkar, Shripad, Nan Li, and Marcus W. Feldman, "High Sex Ratios in China's Future," *Science*, 1995, *267* (5199), 874–876.

Waldron, Ingrid, *Factors Determining the Sex Ratio at Birth*, United Nations, 1998.

Williamson, Nancy E., *Sons or Daughters. A Cross Cultural Survey of Parental Preferences.*, Sage Publications, 1976.

Yamaguchi, Kazuo, "A Formal Theory for Male-Preferring Stopping Rules of Childbearing: Sex Differences in Birth Order and in the Number of Siblings," *Demography*, 1989, *26* (3).

Zeng, Yi, Ping Tu, Baochang Gu, Yi Xu, Bohua Li, and Yongpiing Li, "Causes and Implications of the Recent Increase in the Reported Sex Ratio at Birth in China," *Population and Development Review*, 1993, *19* (2).