Growth and Peace

Essays on economic growth, the causes of war,
AI, and the possibility of a more peaceful future

Ben Garfinkel
Nuffield College



This thesis is submitted for the degree of DPhil in International Relations

Department of Politics and International Relations
University of Oxford
Hilary 2024

Wordcount: 70,172

Abstract

This thesis considers the *Peaceful Future Hypothesis*: the hypothesis that the future will be drastically more peaceful than the past.

I clarify the hypothesis, then describe a two-step approach to judging its likelihood. The first step is to analyse historical trends in the rate of conflict and produce an initial forecast purely on this basis. The second step is to modify this initial forecast, by considering causal mechanisms that could explain past trends — or give us reasons to expect conflict to be more or less common than trend extrapolation would suggest.

After conducting analysis on these lines, **I suggest that a peaceful future is unlikely but not implausible.** (For example, assigning it something like a 15% probability wouldn't strike me as unreasonable.) One finding, which lends the hypothesis some plausibility, is that a noisy downward trend in the rate of interstate conflict seems to have started in the mid-20th-century. If this trend is real and lasts long enough, then the world would be very peaceful by the end of the present century. Economic growth may be *one factor* driving this trend: growth seems to have a strong-but-underrecognized pacifying effect, and many states began experiencing substantial growth around the right time. If growth is actually one driver of the trend, then that would be reassuring, since the world may become much wealthier over the next century.

However, we can't say anything with confidence. Historical trends and causal relationships are hard to reliably identify. Growth is far from the only factor that matters. Growth's marginal effect will probably fall. AI might upend everything. And so on.

Along the way, I make various secondary academic contributions. If nothing else, grappling with the hypothesis turns out to be a good way to grapple with fundamental questions in International Relations.

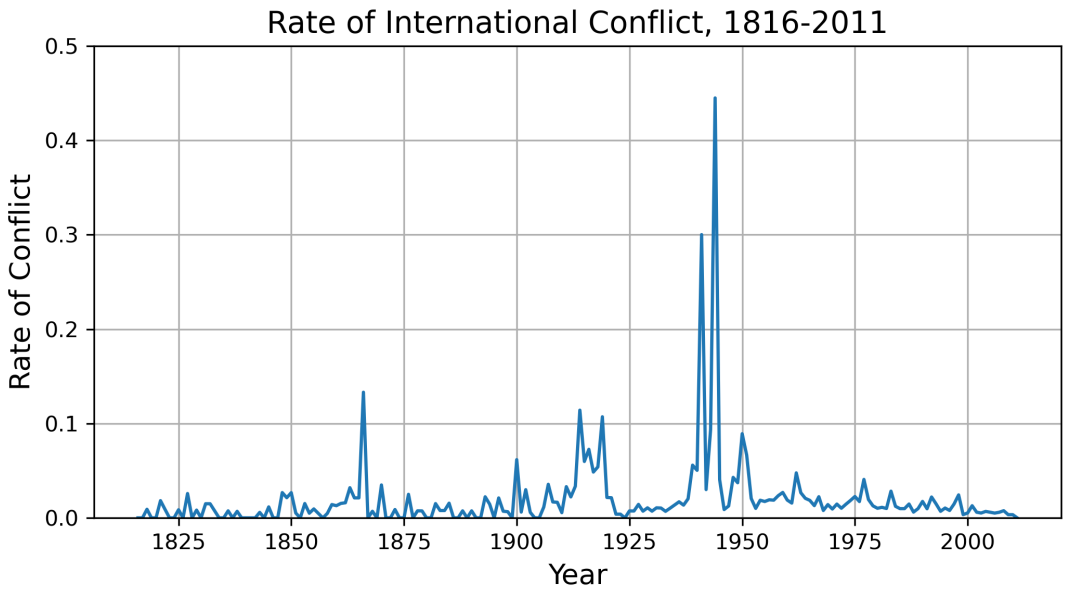

Figure 1: This figure shows the rate of interstate conflict between 1816 and 2011, as implied by the Correlates of War Militarized Interstate Disputes dataset (Palmer, 2015). The rate is defined here as the portion of politically relevant pairs of states that begin to use force against each other in a given year (Braumoeller, 2019). Clearly, any genuine trends and trend breaks in this time series will be difficult to reliably distinguish from random noise.

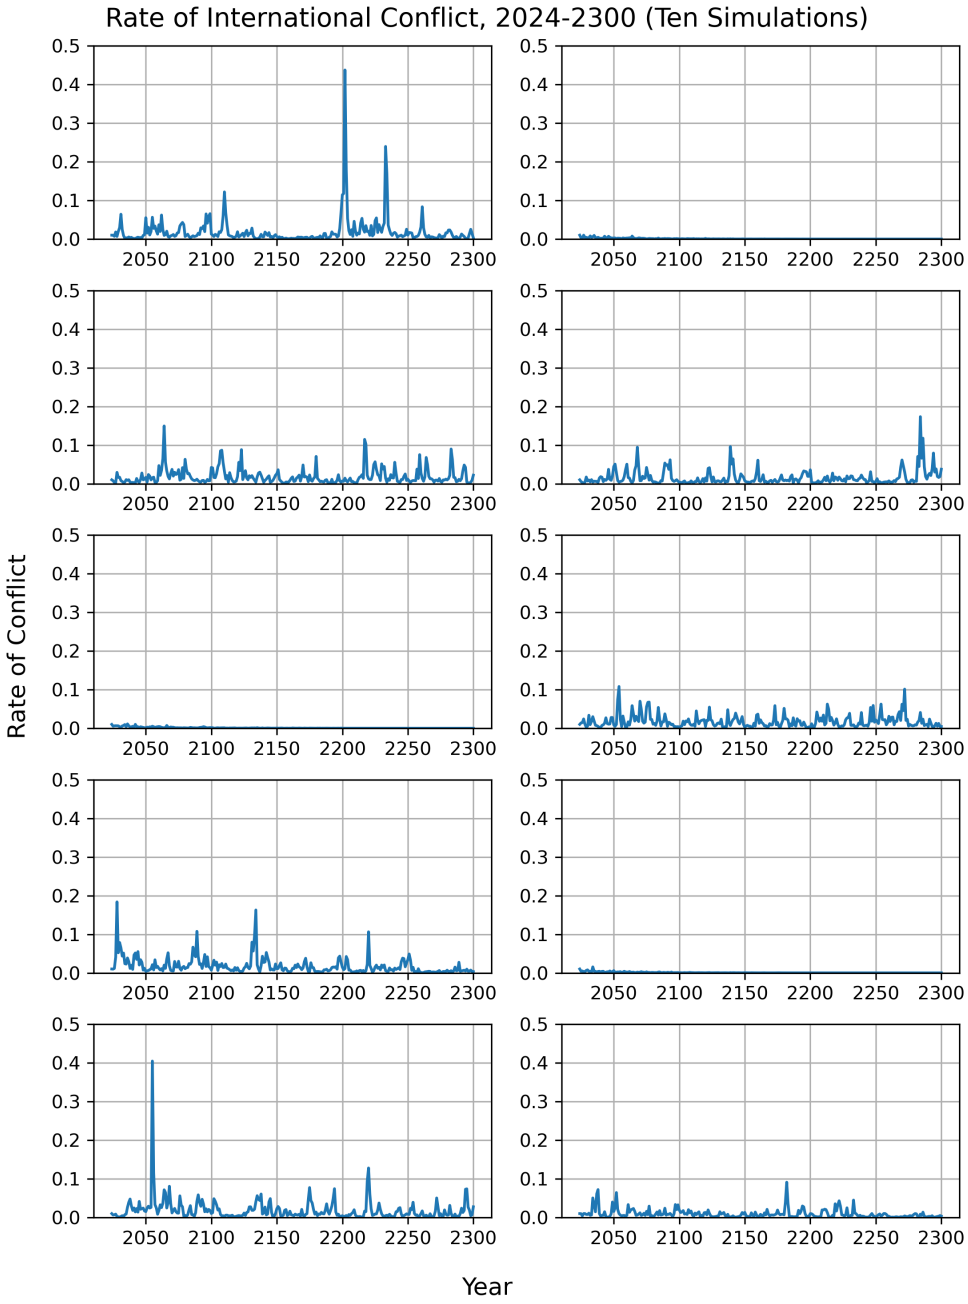

Figure 2: This figure presents ten simulations of the rate of international conflict between 2024 and 2300. They were generated by fitting a model to the historical data and then extrapolating it forward, under the simplifying assumption that the current trend will not change. (Chapter 3 discusses the model.) Both inherent randomness and parameter uncertainty — in particular, uncertainty about what the current trend actually *is* — imply that a wide range of futures could be consistent with historical data. Most of the peaceful futures observed rely on the assumption that a steep trend toward peacefulness began in the mid-20th century.

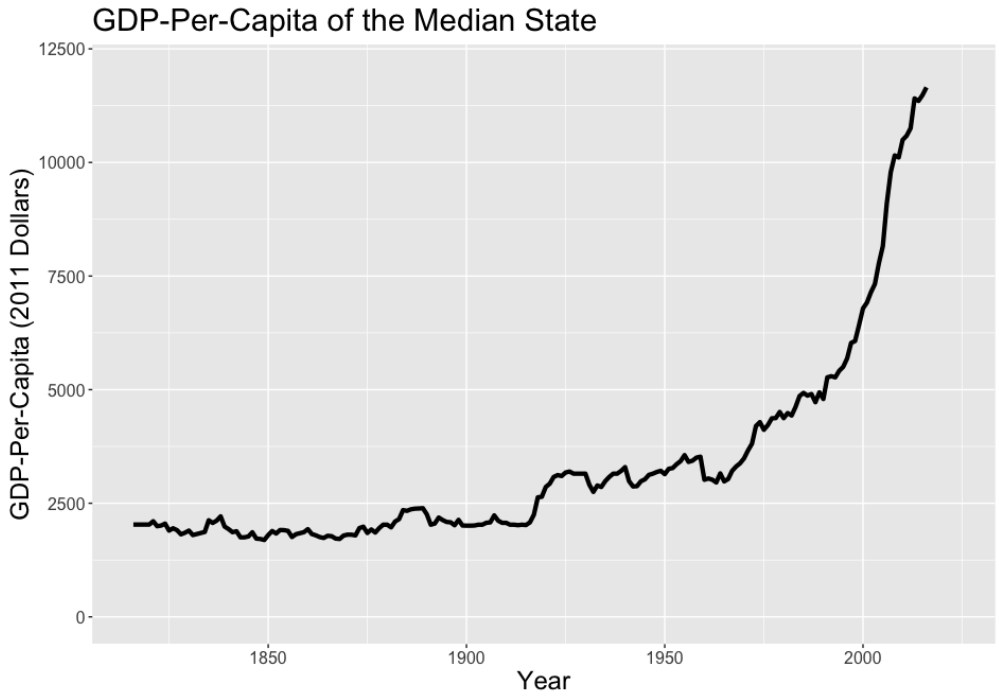

Figure 3: This plot shows the estimated GDP-per-capita of the median state, among states included in the Correlates of War dataset (Bolt and van Zanden, 2020). Widespread economic growth emerged only in the mid-20th century, around the same time that a steep trend toward peace *may* have begun. If growth has a pacifying effect, then it could help explain a recent trend toward peace and thereby increase our confidence that it is real. Continued growth would then also offer a reason for optimism that the trend will be sustained.

Table 1: This regression table — which will be fully explicated in Chapter 5 — reports the results of several logistic regressions. The results suggest that when pairs of states grow more wealthy ("AvgLogGDPPC"), they become much less likely to fight each other. This statistical link between growth and peace has been underappreciated, mainly because it only consistently appears when fixed traits of pairs of states ("dyadic fixed effects") are controlled for. Since there are multiple plausible mechanisms by which growth could support peace, it ultimately seems reasonable to think that growth probably *does* support peace. This conclusion is at least moderately good news.

	Dependent variable:				
	Militarized Interstate Dispute				
	(1)	(2)	(3)	(4)	(5)
AvgLogGDPPC	-0.219***	-0.397****	-0.369***	-0.418***	-0.459***
	(0.069)	(0.118)	(0.126)	(0.141)	(0.148)
AvgDemocracyLevel				-0.077****	-0.081****
				(0.012)	(0.011)
AvgEconDependence				0.000	0.000*
				(0.000)	(0.000)
BothHaveNukes					-0.683**
					(0.270)
EitherHasNuke					-0.562***
					(0.182)
MajorPower					0.672**
					(0.276)
CapabilityRatio					-0.105
					(0.089)
InterestSimilarity					-0.392
					(0.366)
Ally					0.036
					(0.308)
Dyadic Fixed Effects	Included	Included	Included	Included	Included
Time Controls	Not Included	Continuous	Year Dummy	Continuous	Continuous
PeaceYears	Not Included	Included	Included	Included	Included
Geographic Controls	Not Included	Included	Included	Included	Included
Observations	66,154	66,154	66,154	39,623	38,973
Earliest Year	1817	1817	1817	1871	1871
Note:					*p<0.1; **p<0.05; ***p<0.01; ****p<0.001

Acknowledgements

This thesis would not have been possible without the support of my advisors, Duncan Snidal and Allan Dafoe. Their faith in me has been invaluable, and I am hugely grateful for their guidance, openness, and patience. Especially their patience.

Allan, in particular, was my first academic mentor. No one else has had a larger impact on my intellectual development. Even though I have failed to give him an update on this thesis for ages, his fingerprints are all over it.

Beyond providing valuable input on my research, Duncan played a critical role in supporting my application to the DPhil program. At the time, my academic background was in entirely unrelated fields, and I had only a single idiosyncratic International Relations publication. He was willing to take a bet on me, then support my inevitable decision to write an idiosyncratic International Relations thesis.

This thesis also would not have been possible without my parents. They not only caused me to exist, but also — as my mother sometimes puts it — raised me right.

A core paper in this thesis also would not have been possible without the late Bear Braumoeller. While I never had the privilege of meeting him, I am grateful to him for *doing good work on questions that matter*. As academics, that is what we all should aspire to.

Realistically, this thesis probably *would* have been possible without my colleagues and friends. In fact, since I would have had more free time if I were unemployed and alone, I surely would have finished my thesis sooner. However, I believe they deserve to be forgiven.

Contents

Acknowledgements **6**

1 Introduction **13**

1.1 Optimistic Visions 13

1.2 What I Mean By a Peaceful Future 15

1.3 Focusing on Growth 16

1.3.1 What I Mean by Growth 16

1.3.2 Why Focus on Growth? 17

1.4 Outline of Thesis 19

1.5 Relationship to Existing Literature 22

1.6 A Note On Prerequisites 23

2 Peace Optimism, Past and Present **25**

2.1 Abstract 25

2.2 Ancient and Early Modern Thinking 25

2.3 Peace Optimism in the Early Nineteenth Century 27

2.4 "Norman Angellism" and Other Pre-WWI Thought 30

2.5 Interwar Thought 33

2.6 Early Cold War Thought 34

2.7 The Return of Peace Optimism 36

2.8 The New Optimists 37

2.9 The New Pessimists 39
2.10 Conclusion 41

3 Will the Future Be Peaceful? How Would We Know? 44

3.1 Abstract 44
3.2 Introduction 45
3.2.1 The Peaceful Future Hypothesis 45
3.2.2 Overview of Paper 47
3.3 Clarifying the Peaceful Future Hypothesis 49
3.4 Forecasting and Political Science 51
3.4.1 In Defense of Forecasting 52
3.4.2 Forecasting: What Works 54
3.4.3 What Intellectual Progress Looks Like 58
3.5 How Can We Forecast the Future of Conflict? 59
3.5.1 A Two Stage Process 59
3.5.2 Plan for the Remainder of the Paper 60
3.6 Trend Extrapolation: Forming an Initial "Outside View" Forecast 61
3.6.1 Basic Approach 61
3.6.2 Choosing a Metric for the Rate of Conflict 61
3.6.3 Choosing a Dataset 69
3.6.4 Refining Our Forecasting Question 73
3.6.5 Identifying and Extrapolating Trends 74
3.7 Various Complications: Modifying Our Initial "Outside View" Forecast . . . 88
3.7.1 Structural Breaks 90
3.7.2 Data Uncertainty 98
3.7.3 Model Uncertainty 100
3.8 Summary: Where Our Forecast Stands 102
3.8.1 How to Interpret Our Outside View Forecast 102

3.8.2 How Our Choice of Question Has Mattered 104
3.9 Next Steps: Incorporating "Inside Views" 105
3.10 Conclusion 107
3.11 Appendices 108
3.11.1 Appendix 1: Other Datasets 108
3.11.2 Appendix 2: Accounting for Data Uncertainty 114

4 Growth and Civil War **121**
4.1 Abstract 121
4.2 Introduction 121
4.3 Background 124
4.3.1 Prosperity and Peace: The Country-Level Statistical Relationship . . 124
4.3.2 Possible Mechanisms Linking Growth to Peace 128
4.4 The Diminishing Returns Mechanism 130
4.4.1 Risk-Taking and Financial Incentives: The Core Model 131
4.4.2 Accounting for Additional Incentives: A More General Model 137
4.4.3 Does Prosperity Actually Matter? 140
4.4.4 Implications 143
4.5 Other Relevant Empirical Evidence 144
4.5.1 Risk Avoidance in Other Domains 145
4.5.2 Insurgent Responsiveness to Shifts in Economic Incentives 146
4.5.3 Insurgent Income Premiums 147
4.5.4 Summary 147
4.6 Conclusion 148
4.6.1 Recapitulation 148
4.6.2 What About Interstate War? 148

5 Growth and Interstate War **150**

5.1 Abstract . . . 150

5.2 Introduction . . . 151

5.3 Previous Work . . . 155

5.4 Growth and Interstate War: The Diminishing Returns Mechanism . . . 157

5.4.1 A Bargaining Model . . . 157

5.4.2 Further Considerations . . . 167

5.5 Other Effects of Growth . . . 173

5.5.1 Changes in the Mobility and Fragility of Wealth . . . 173

5.5.2 Changes in Economic Interdependence . . . 174

5.5.3 Changes in the Prevalence of Democracy . . . 175

5.5.4 Changes in Culture . . . 178

5.5.5 Changes in Young Men's Access to Casual Sex . . . 179

5.5.6 Changes in Demographic Structure . . . 180

5.5.7 Changes in the Offense-Defense Balance . . . 181

5.5.8 Changes in the Information States Have About Each Other . . . 182

5.5.9 Should We Expect the Marginal Effect of Growth Decline? . . . 183

5.5.10 Summary . . . 184

5.6 Growth and the Correlates of Interstate War . . . 184

5.6.1 Previous Studies . . . 184

5.6.2 The Importance of Fixed Effects . . . 186

5.6.3 Research Design . . . 190

5.6.4 Results . . . 195

5.6.5 Discussion . . . 196

5.7 Conclusion . . . 201

5.8 Appendices . . . 202

5.8.1 Calculating the Approximte Effect of a Doubling in GDPPC . . . 202

5.8.2 Rerunning the Regressions with Older GDP Data 203

6 Adjusting Our Initial Forecast 206

6.1 Abstract 206

6.2 Background 207

6.2.1 Our Outside View Forecast 207

6.2.2 Our Understanding of the Effect of Growth 207

6.3 Growth and the History of Conflict 214

6.3.1 Can Growth Explain the (Apparent) Downward Trend in Conflict? . 214

6.3.2 Other Interpretations 216

6.3.3 Summary 220

6.4 Growth and the Future of Conflict 220

6.4.1 Forecasting Growth 221

6.4.2 Implications of Different Growth Scenarios 229

6.4.3 A Conditional Forecast 230

6.5 Conclusion: A New Overall Forecast 231

6.6 Appendix: Complications From Uneven Growth 233

7 Conclusion 235

7.1 Abstract 235

7.2 The Thesis So Far 235

7.2.1 The Central Conclusion 235

7.2.2 How We Arrived at This Conclusion 236

7.3 Academic Contributions 237

7.4 Directions for Further Work 238

7.5 Closing Thoughts: Fundamental Questions in International Relations 240

Bibliography 242

A The Risk of Conflict in Our Lifetimes **253**

A.1 Are We All Going to Die? 253

A.2 Some Cursory Analysis 255

B Advanced Artificial Intelligence and the Future of Conflict **259**

B.1 The Possibility of Revolutionary Change 259

B.2 Potential Influence Pathways 261

B.2.1 Changes in Fundamental Rationalist Barriers to Peace 261

B.2.2 Changes in Economic Conditions 272

B.2.3 Changes in the Prevalence of Democracy 276

B.2.4 Changes in the Relationship Between Offense and Defense 278

B.2.5 Changes in International Order 282

B.2.6 Changes in Culture and Values 284

B.3 Conclusion 289

Chapter 1

Introduction

1.1 Optimistic Visions

In August 1849, nearly one thousand delegates from Britain, France, the United States, Belgium, Italy, and Germany gathered in Paris to discuss the abolition of war (Gilpin, 1849). The delegates to this meeting of the recently formed International Peace Congress included some of the most prominent intellectuals of the time, including Victor Hugo, Richard Cobden, and Frédéric Bastiat. On the first night, Hugo, serving as the conference's president, presented a strikingly optimistic outlook. The complete elimination of war was not only possible, he declared, but ultimately inevitable. The task of the gathered delegates was merely to accelerate a "tendency toward unity" and a "softening of manners" that was already underway:

> *Gentlemen, this sacred idea, universal peace...I ask you, is it practicable? Can it be realized? Many practical men, many public men grown old in the management of affairs, answer in the negative. But I answer with you, and I answer without hesitation, Yes!...I go still further. I do not merely say it is capable of being put into practice, but I add that it is inevitable, and that its execution is only a question of*

time, and may be hastened or retarded.... A day will come when a cannon will be a museum-piece, as instruments of torture are today. And we will be amazed to think that these things once existed!

Hugo's speech belonged to an already well-established intellectual tradition of *peace optimism*. Distinct from pacifism, the normative view that war is unethical, peace optimism is the positive view that war is likely to become dramatically and lastingly less common.[1] Some of its earliest proponents were prominent Enlightenment philosophers such as Immanuel Kant and Jeremy Bentham. Since this time, peace optimism has repeatedly disappeared and reappeared in mainstream political thought.

Peace optimism has again reappeared over the past few decades. Several prominent academics, most notably Steven Pinker, have helped draw attention to an apparent long-run decline in interstate war (Keegan, 1993; Mueller, 2004; Pinker, 2011; Goldstein, 2011; Morris, 2014; Gat, 2017). To explain the apparent decline, these writers have also in many cases returned to the arguments of early peace optimists such as Kant. Although Pinker avoids explicitly predicting the disappearance of war, others have been less circumspect. For instance, writing in *International Security*, Michael Mousseau writes that "the world is now in the endgame of a five-century-long trajectory toward permanent peace and prosperity" (Mousseau, 2019). While peace optimism may again be receding, due to recent conflicts such as the Russo-Ukrainian war, history suggests that we should await its return.

This thesis is motivated by a desire to understand whether, in the long run, a more peaceful future is actually plausible.



Figure 1.1: British casualties of the Crimean War, which broke out four years after Hugo's speech. The Crimean War interrupted a long period of relative peace inside Europe and temporarily dampened hopes for the disappearance of war. Over the past two centuries, peace optimism has risen and fallen across multiple waves.

1.2 What I Mean By a Peaceful Future

The central preoccupation of this thesis is the possibility of what some authors have called a "permanent peace" or a "perpetual peace." I will instead call it a "peaceful future." While there are many possible standards that could be used to define the concept, some extremely strict, the standards I have in mind are at least somewhat permissive.[2] I do not have in mind a world without conflict, a demilitarized world, or a fundamentally harmonious world. What I have in mind is the following hypothesis:

[1] "Peace optimism" is not an established term. However, to my knowledge, there is no commonly used term that captures precisely this concept. The closest term may be "pacificism," which is sometimes used to describe the views of individuals who strive to reduce the prevalence of war but do not fully endorse pacifism (Ceadel, 1996). Even this term, though, primarily describes a normative position.

[2] It is unusual for authors discussing "perpetual peace" or "permanent peace" to define the term. Particularly in the case of eighteenth and early nineteenth century writing, this ambiguity has led to significant interpretive debates (Aksu, 2008). For instance, early writers may often have used the term to refer specifically to peace among European states and largely disregarded wars waged outside European borders.

> **Peaceful Future Hypothesis:** There will be (or has already started to to be) an extremely large and long-lasting decline in the rate of interstate conflict, relative to its historical level.

Of course, this definition is still quite vague. In fact, one of the key objectives of the third chapter of this thesis will be to make the hypothesis more precise. Still, even this vague definition can serve as an entry point for discussion.

1.3 Focusing on Growth

There are many possible stories about how a peaceful future could be reached. Historically, optimists about peace have tended to focus on the potential pacifying effects of the spread of democracy, the expansion of international trade, the development of more effective international institutions, and continued moral and intellectual progress.

Although I will discuss each of these stories, to some extent, I will focus most strongly on the potential effect of *economic growth*.

1.3.1 What I Mean by Growth

What I mean by "economic growth" is this:

> **Economic growth:** an increase in a society's production of valued goods and services, as well as any qualitative economic changes (such as industrialization and the adoption of new technologies) that make this increase in output possible.[3]

Within this thesis, the term "economic growth" can also be treated as roughly synonymous with the term "economic development."

[3]Economic growth can also be framed in terms of a rise in "income," rather than a rise in "output," since output and income go together. The more valuable goods and services are produced (output), the more valuable goods and services are available to be consumed (income). A society that produces very little stuff is also a society that is very poor.

Economists draw a distinction between growth in output-per-person (e.g. GDP-per-capita) and growth in total output (e.g. GDP). When I use the term "economic growth," you should generally assume that I am referring to growth in output-per-person. However, so long as total output grows more quickly than population levels, these two forms of growth go together.[4]

As a last piece of clarification, I will be primarily interested in the *lasting effects* of growth rather than the effects of growth *while it is occurring*. In other words, I will be looking at the implications of countries reaching increasingly high *levels* of output rather than the implications of countries having high output *growth rates*.

1.3.2 Why Focus on Growth?

In general, my decision to give special attention to an individual causal factor is dictated by pragmatism. Even in a full-length thesis, it is impossible to *deeply* analyze every factor that might play a role in determining how peaceful the future will be.

There are, then, a few reasons why I have ended up focusing on growth.

First, a priori, growth seems like a natural candidate for a factor that could have a *large* impact on the rate of international conflict. Growth tends to dramatically transform countries as they climb up the ladder of economic development. Reflect, for example, on the vast difference between pre-industrial Britain and modern Britain. (See Figure 1.2.) Fairly reliably, as countries grow, poverty gives way to abundance, people move out from the fields and into the cities, land falls in importance as a source of wealth, economic specialization and division of labor expand, and people start living longer and having fewer children, among many other changes. It would not be surprising if at least some of these changes had important implications for the risk of conflict. Perhaps more interestingly, growth is also a natural candidate for a factor that could have a *continuing* impact on the rate of international con-

[4]Up until the past two centuries, however, absolute growth rates were generally too low to allow output-per-person to meaningfully rise. As will be discussed later on this thesis, sustained and large-scale growth in output-per-person is a distinctly modern phenomenon.

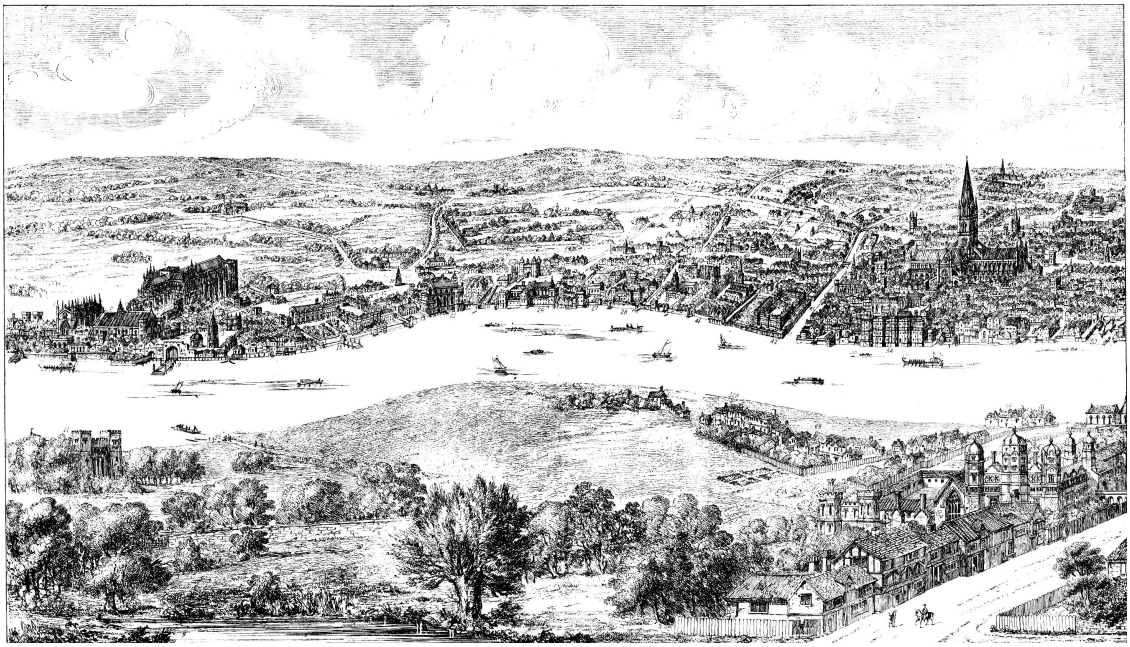

Figure 1.2: A view of London in 1543, roughly two centuries before Britain began to experience sustained economic growth. So far, growth has brought a number of changes. (Engraving by Nathaniel Whittock, based on a drawing by Antony van den Wyngaerde.)

flict. Growth trends within countries are mostly one-directional and surprisingly consistent over long stretches of time.

Second, the potential impact of growth for international conflict is still relatively underexplored. Growth (or "development") has been called "the neglected stepchild of the study of economics and conflict" (Gartzke, 2015). Discussions of the future of conflict discuss growth much less often than they discuss, for instance, the spread of democracy and economic interdependence. Gat (2017), who expresses an optimistic view about the future of conflict and attributes his optimism in large part to the effects of growth, is notable as an exception.

Third, looking ahead, there is some reason to think that growth may accelerate in the coming century. In particular, speculatively, some leading economists have suggested that AI-driven automation could eventually lead growth to accelerate (Aghion et al., 2018; Nordhaus, 2021; Trammell and Korinek, 2023). This possibility gives forward-looking researchers an especially large reason to try to understand the implications of growth for conflict.

1.4 Outline of Thesis

The opening (literature review) chapter of the thesis offers an intellectual history of peace optimism, from the Enlightenment to the present day. It identifies the enduring themes in arguments for and against the Peaceful Future Hypothesis, and it reviews the present state of the debate.

The first paper is entitled: "Will the Future Be Peaceful? How Would We Know?" In this paper, I argue that it is possible to productively study the Peaceful Future Hypothesis and make progress on assessing its plausibility. I start out by operationalizing the hypothesis: **I focus on the possibility that the average rate of conflict, according to one particular metric, will be *at least 90% lower* between 2050 and 2300 than it was between 1850 and 2000.** I then describe a two-stage framework for grappling with the hypothesis. The first stage involves studying historical trends in the rate of conflict and producing an initial forecast purely on this basis. The second stage involves modifying the initial forecast, by considering different causal factors that could explain historical trends or that could give us reasons to expect the rate of conflict to overshoot or undershoot them. I then walk through the stages of producing an initial trend-based forecast. One interesting finding is that the risk of conflict seems to have been trending downward since the mid-20th century, albeit in a noisy and erratic way. Overall, the initial forecast suggests that a peaceful future is unlikely but at least plausible. I give 20% as an example of an initial probability estimate that I would consider reasonable.

The next two papers then move beyond trend analysis and consider a specific causal factor that could play an important role in determining the future rate of conflict. Specifically, in these papers, I try to clarify the effect of *economic growth* on the risk of conflict. Since economic growth is such a transformative phenomenon, and since it will probably continue for many more decades, it is a particularly natural factor to consider.

I start out (in "Growth and Civil War") by considering the effect of growth on the risk of civil war. My reasons for initially focusing on civil war are that: (a) this case is in

some ways simpler than the interstate war case and (b) scholars already generally believe that growth has a pacifying effect in this case. I review the state of empirical evidence for a causal relationship between growth and peace within countries and review the most popular mechanistic explanations for this relationship. I then argue that one particularly blunt explanation has been prematurely dismissed. Specifically, as people become wealthier, they experience strongly diminishing returns on further wealth; as a result, it becomes harder for financial rewards to motivate them to risk their lives or freedom in conflicts. This Diminishing Returns Mechanism may play an important role in linking economic growth to domestic peace.

I then (in "Growth and Interstate War") turn to directly consider the case of interstate conflict. Unlike in the case of civil war, International Relations scholars do not already (for the most part) believe that growth reduces the risk of interstate conflict. However, I argue that growth probably *does* make states less likely to fight each other. I note that we should expect the Diminishing Returns Mechanism to also apply in the case of interstate conflict. I then survey many additional causal mechanisms that other scholars have put forward, which may allow economic growth to influence the rate of interstate conflict. Finally, I report that there is a strong and surprisingly underrecognized statistical link between growth and interstate peace. As pairs of states grow wealthier, they become much less likely to fight each other: in particular, *if two states become twice as wealthy, the chance of them fighting seemingly falls by about a third.* (For reference, the median state has become about four times as wealthy over the past century.) Ultimately, available arguments and evidence suggest that growth pushes in the direction of international peace. At the same time, though, it is worth keeping in mind that the effect of growth may decline as states become wealthier — and that there are also pathways by which growth can actually increase the risk of conflict.

A synthesis chapter returns to our initial trend-based forecast of the future of conflict, in light of our newly improved understanding of the effect of economic growth. I argue that economic growth is one possible, partial explanation for the downward trend in the rate of

interstate conflict that seems to have begun in the mid-20th century: most of the world's states also began to experience substantial economic growth around this time. Growth also seems likely to continue for many more decades and, due to progress in artificial intelligence, may even accelerate. We therefore have an additional reason, beyond pure trend extrapolation, to consider a peaceful future *plausible*. Nonetheless, I conclude, we would still not be justified in considering a peaceful future *likely*. The analysis in this chapter actually leads me to adjust the initial probability estimate slightly downward, arriving at a final figure of 15%. This specific number should not be taken too seriously, since somewhat higher or lower probability estimates would also be about equally defensible. However, I believe it would that it would at least be difficult to defend very low estimates (e.g. 1%) or very high ones (e.g. 60%). My sense is that the *spectrum of reasonable disagreement* about the likelihood of a peaceful future — at least given the set of evidence and considerations covered by this thesis — is concentrated largely around the low double-digits.

The conclusion chapter summarizes the preceding chapters, notes some directions for future work, and closes with a few brief additional reflections.

...

The thesis also has two appendices.

The first appendix briefly discusses the distinct-but-related question of how likely a major war is to occur within the next few decades. I don't offer a forecast, but I do sketch a high-level approach to producing this forecast and explain how the analysis in this thesis should feed into it.

The second appendix then offers a deep dive into other possible implications of progress in AI, beyond its potential impact on economic growth. A key take-away is: *There are actually several plausible pathways by which AI could heavily alter the risk of conflict.* Unsurprisingly, it is difficult to form a view about what the net effect will be. Nonetheless, while I have not yet attempted to integrate this AI-focused appendix's analysis into my over-

all forecast, my *tentative* view is that it probably does overall offer grounds for additional optimism.

1.5 Relationship to Existing Literature

The chapters in this thesis draw on and contribute to a number of distinct literatures.

The first paper, "Will the Future Be Peaceful? How Would We Know?", builds on an existing literature on long-run trends in conflict and their implications (e.g. Pinker, 2011; Cirillo and Taleb, 2016; Mearsheimer, 2018; Braumoeller, 2019). I build most directly on Braumoeller (2019). I go substantially beyond this work, though, in part by drawing in frameworks and modelling tools from literatures such as time series analysis, political forecasting (e.g. Tetlock, 2016), and Bayesian model selection (e.g. Wasserman, 2000). This paper is also unique in that it ventures to put forward a probabilistic forecast of the future rate of conflict.

The second paper, "Growth and Civil War," draws on a large existing literature on economic correlates and causes of civil war (e.g. Fearon and Laitin, 2003; Blattman, 2010; Ray, 2017). Its main contribution is to develop and argue for a previously neglected explanation for the "puzzle" that wealthier states are much less likely to experience civil wars. It also, in a more indirect way, contributes to the "greed or grievance" literature on combatant motives (e.g. Collier and Hoeffler, 2004; Hoeffler, 2011; Keen, 2012); it develops a rational actor model that allows financial and non-financial motives to co-exist and clarifies how they interact. The modelling work in this paper draws fairly directly on the economics literature on income and risk aversion; it takes particular inspiration from a rational actor model by Hall and Jones (2007), which was developed in part to explain why people in wealthier countries spend larger portions of their incomes on healthcare.

The third paper, "Growth and Interstate War," is mainly situated within the Capitalist Peace literature (Gartzke, 2007; Schneider, 2017). This literature explores the potential

pacifying effects of aspects of the modern economy, especially international trade. One contribution of the paper is to clarify and argue for one potential causal mechanism linking economic growth to peace, which seems to have only received a brief mention in the literature (Gat, 2017). A second contribution is to clarify and synthesize many of the other causal mechanisms described in the Capitalist Peace literature, by incorporating them into a unified model that is grounded in bargaining theory (Fearon, 1995). A third contribution is to offer statistical evidence of a strong link between growth and interstate peace. My statistical findings diverge from those of previous studies, most of which have failed to report a link, mainly because I rely on fixed-effects models that (I argue) we should expect to more effectively control for potential confounders.

Overall, in addition to Braumoeller (2019), Gat (2017) is the previous work that is mostly closely related to this thesis. I would recommend his book to any reader who is interested in the subject of this thesis. Gat also argues that the rate of war is declining, argues that economic growth may be one factor driving this apparent decline, and notes the Diminishing Returns Mechanism as one mechanism that may be at play. This thesis takes a fairly different approach from Gat, though, in part through its extensive use of formal methods. Distinctive aspects of this thesis include: its statistical analysis of historical trends, its use of rational actor models to formalize theoretical arguments about the effect of growth, and its effort to produce an explicit forecast.

1.6 A Note On Prerequisites

All of the papers that make up this thesis assume that the reader is familiar with the basics of regression analysis and the basics of algebra. In general, some level of comfort with math is expected.

However, I have tried not to assume much prior knowledge beyond this. I have tried not to assume, for example, that the reader has a substantial background in International

Relations. This means that I will sometimes explain points that would normally be taken for granted. In some places, the thesis takes on a kind of textbook-like flavor.

One reason I do not assume a strong International Relations background is that I can at least *imagine* people from a few different fields finding this thesis interesting. It is unusually multidisciplinary and addresses an unusually "big picture" topic. Of course, the typical PhD thesis is read by effectively no one beyond one's advisors and examiners. But if someone else *does* want to read this one, they might not be an International Relations scholar.

Chapter 2

Peace Optimism, Past and Present

2.1 Abstract

In this (literature review) chapter, I review past and present thought regarding the possibility of a peaceful future. I show that peace optimism has now risen and fallen in popularity several times, but that its core themes have remained surprisingly consistent. I close by discussing lessons from this history and promising directions for further research on the plausibility of a peaceful future. Some of these directions will be explored by the papers that make up this thesis.

2.2 Ancient and Early Modern Thinking

Hope for a permanent end to war may be as old as war itself. Our oldest surviving records of this hope, however, date back to the Axial Age. The Hebrew Bible famously prophesizes that God will eliminate the need for war, by instituting a kind of international arbitration process:

> *He shall judge between the nations, and shall arbitrate for many peoples; they shall beat their swords into plowshares, and their spears into pruning hooks; nation shall*

not lift up sword against nation, neither shall they learn war any more.

At approximately the same time, a number of Chinese philosophers were responding to the violence of the Warring States period with discussions of the roots of war (Yates, 2016). Mozi, the founder of Mohism, discussed the possibility of a peace founded on "universal love." Less prosaically, Mohists also assigned great significance to what would now be called the "offense-defense balance": they worked to spread knowledge and technology that would differentially benefit states under siege.

Two thousand years later, an especially sophisticated and secular strand of thinking on war and peace spread across Europe. The European Enlightenment had ignited widespread optimism about the possibility of intellectual progress. A number of philosophers therefore began to conceptualize war as a puzzle to solve. Immanuel Kant, Jean-Jacques Rousseau, and Jeremy Bentham all contributed to a literature arguing that certain political reforms could enable "perpetual peace" (Aksu, 2008).

Bentham's "Plan for an Universal and Perpetual Peace," written shortly before the French Revolution, is perhaps the most sophisticated example of this literature (Bentham, 1927). Bentham's central claim is that, at least for people below the elite strata of each country, the costs of wars tend to outweigh their benefits. To eliminate war, then, it is of greatest importance to increase the power of common people and to increase awareness of war's costs. It will also be important to reduce the number of opportunities for misunderstandings and conflicts of interest to arise between states and to develop alternative methods of dispute resolution. On these theoretical grounds, Bentham argues for a program to spread inclusive domestic political institutions, educate the public about matters of war and peace, increase government transparency, support free trade, dissolve European colonies, avoid security commitments, develop a body of unambiguous international law, and establish international arbitration bodies.

There is significant overlap between these prescriptions and the prescriptions in Kant's

later and much more well-known essay "Perpetual Peace: A Philosophical Sketch" (Kant, 1983). In addition to several "preliminary articles" of secondary importance, Kant proposed three "definite articles" that would produce perpetual peace if adopted by all states. States should adopt republican forms of government, establish an international legal institution (a "league of peace"), and extend "universal hospitality" toward one another's citizens. This final definite article is often interpreted as implying free trade.

Inclusive domestic political institutions ("democracy"), international institutions, and trade would later come to be known as the "Kantian tripod" (Russett et al., 1998). From the late eighteenth century to the present day, the majority of writers who have raised the possibility of perpetual peace have justified their proposals or predictions by evoking at least one of the triangle's corners. Moral and intellectual progress is the only other candidate pacifying force that has been cited nearly as often.

Both Bentham and Kant were outright optimists about the future of war. Bentham expressed the view that the elimination of interstate war was not only possible, but very likely to occur at some point in the future (Conway, 1989). Kant's optimism, founded on a teleological understanding of human history, was even bolder. In a supplement to "Perpetual Peace," he famously claims that perpetual peace is nothing less than the destiny of humanity:

> *The guarantee of perpetual peace is provided by nothing less than that great artist called Nature. In Nature's mechanical progress we see that her aim is to produce a harmony among men, against their will and actually through their discord. As a necessity working according to laws we do not know, we submit to it as Destiny.*

2.3 Peace Optimism in the Early Nineteenth Century

An international peace movement first took shape in the 1810s, then grew over the next several decades (Laity, 2002; Cortright, 2008). Dozens of "peace societies" sprung up

first in Britain and the United States, then later in continental Europe. In Britain, the leading society was the Quaker-led British Society for the Promotion of Permanent and Universal Peace. Although these societies were initially religious in nature, and focused particularly on moral persuasion as a path to peace, the movement quickly attracted more secular and pragmatic members.

In the first half of the nineteenth century, the high water mark for the movement was arguably the 1849 International Peace Conference. Victor Hugo's opening address, quoted in the introduction to this thesis, expressed an unqualified optimism about the future of war. The peace movement was quite clearly gaining in momentum. Furthermore, in Hugo's view, war now was bound to disappear no matter what the movement chose to do. Their purpose was simply to "hasten" the "inevitable" arrival of perpetual peace.

A few key pieces of historical context help to explain why peace optimism, although certainly not a majority view, became particularly common in the first half of the nineteenth century. First, initially, the protracted bloodshed of the Napoleonic Wars created a strong feeling of war weariness in many European states. A desire to avoid repeating these conflicts was commonly felt. On the other side of the Atlantic, disgust with the outcome of the War of 1812 combined with a Christian religious revival to create a similarly favorable environment for peace movements.

Second, the long nineteenth century was in many ways a period of unprecedented human progress. Scientists and engineers developed a number of technological marvels, such as the telegraph and the smallpox vaccine, which lent credence to the idea that careful study could produce solutions to ancient problems. Perhaps more importantly, the American and French Revolutions had shown that radical institutional change was possible. The subsequent successes of the British abolition and voting reform movements, which were intertwined with the peace movement, provided important encouragement. If a practice as ancient and odious as slavery could be placed on a path toward elimination, then, it stood to reason, perhaps war could be placed on this path too.

Third, in Europe, the Napoleonic Wars were followed by a period of unprecedented peace. As decade after decade passed without another great power war, there was increasing empirical evidence that European states were finally learning to avoid conflict with one another. Shortly after the Crimean War interrupted this long period of peace, the prominent British historian Henry Thomas Buckle reflected on the abnormality of the past four decades (Buckle, 1861):

> *That this barbarous pursuit [war] is, in the progress of society, steadily declining, must be evident, even to the most hasty reader of European history. If we compare one country with another, we shall find that for a very long period wars have been becoming less frequent; and now so clearly is the movement marked, that, until the late commencement of hostilities, we had remained at peace for nearly forty years: a circumstance unparalleled, not only in our own country, but also in the annals of every other country which has been important enough to play a leading part in the affairs of the world.*

Finally, a number of trends that emerged in the nineteenth century suggested plausible *mechanisms* for the observed decline in war. For instance, like many others authors, Buckle emphasized the declining power of elites and the growing interconnectedness of European states (roughly corresponding to two corners of the Kantian triangle) as key factors explaining the long peace. Since he believed that these trends were unlikely to be reversed, it naturally followed that war's decline would continue.

Overall, in the four decades that followed the Napoleonic Wars, several peculiar conditions came together to foster peace optimism.

2.4 "Norman Angellism" and Other Pre-WWI Thought

The 1850s and 1860s were a period of the decline for peace movements on both sides of the Atlantic. In the United States, where the peace movement was also closely intertwined with the abolitionist movement, many peace activists became vigorous supporters of the Union cause. In Europe, the Crimean War and several mid-century wars of national unification had a similar dampening effect on peace advocacy. Advocacy for the elimination of war was difficult to mix with passionate support of particular wars.

As illustrated by Buckle's quote above, some intellectuals continued to express optimistic views in this period. Nonetheless, the case for peace optimism had also grown weaker. It was no longer so obvious that war was disappearing. Instead, if anything, war seemed to be on the rise. Furthermore, following the failed revolutions that swept Europe in 1848, the long-run success of parallel reformist movements was now more in doubt. Some of the key justifications for peace optimism in the first half of the nineteenth century no longer held.

Gradually, however, both the peace movement and intellectual interest in perpetual peace rose to new heights (Cortright, 2008). By the first decade of the twentieth century, the peace movement had become truly mainstream. Prominent political figures such as Teddy Roosevelt now publicly advocated for the creation of a league of nations with the power to enforce international agreements. Tsar Nicholas II had organized a major international peace conference at the Hague. Peace societies and discussion groups bloomed across Europe and the United States. Writing after the outbreak of the First World War, the political scientist Stephen Leacock attempted to characterize and justify the sense of hopefulness that he and many other intellectuals had felt (Leacock, 1916). It is worth quoting his essay at length:

> *[I]n earlier ages the war of tribe against tribe and the fight of man against man seemed a necessity of human destiny. Permanent peace could only be a spiritual aspiration. . . . Then came the wonderful age of the nineteenth century, differing utterly from every period of history that had preceded it. Since the coming of the steamboat and the steam*

engine, about a hundred years ago, the world has witnessed a mechanical and industrial progress that has transformed every phase of human life. Man seemed triumphant over nature. The new power of machinery multiplied a hundred fold the means of sustenance. The new possibilities of communication made the world's thought one. The institution of democratic government seemed to supply a mechanism of common action and universal agreement. The interchange of commerce appeared to connect the welfare of each with the welfare of all. War began to appear a horrible and impossible thing, a mere relic of the past, something that could never again, in the light of common sense and common thought, devastate a world of peace and plenty. Prophets of peace abounded. It was freely said that the days of war were ended. "It is quite out of the question," wrote a distinguished English advocate of peace as recently as 1913, "that a single soldier of all the enormous German army will ever live to see a gun fired in war." This and similar expressions of opinion were part of the common stock in trade at the opening of the twentieth century.

Interestingly, as this passage may suggest, the core themes in early twentieth century peace optimism were very similar to the core themes in early nineteenth century peace optimism. The Kantian triangle of democracy, trade, and international institutions was still central. Rapid scientific and social progress in other domains still provided important inspiration. A long period of relative peace in Western Europe also made it easy for some intellectuals to believe that they had already observed a robust decline in the likelihood of interstate war in this part of the world.

However, the ideas of peace optimists had in fact grown more sophisticated. The most important intellectual milestone in this area was the publication of Norman Angell's highly influential book *The Great Illusion* (Angell, 1913). Angell developed, in unprecedented detail, the argument that trade raises the costs of war by fostering interdependence. He also proposed a possibly novel argument that the profitability of conquest declines as states

undergo economic development. Since wealth in the developed world is "founded upon credit and commercial contract," and since this form of wealth is "undermined" by attempts to confiscate it, attempts to confiscate wealth should become less efficient over time. In addition, because development tends to support the emergence of a unified national identity within a state's borders, the likelihood of intense resistance to conquest should grow too.

Angell's book primarily argued that war is becoming increasingly *irrational*, rather than arguing that it was becoming increasingly *unlikely*. However, it did also clearly raise the prospect of a permanent European peace. Angell acknowledged that belligerence, pride, and irrational belief in the profitability of war were lingering barriers to peace, but suggested that humanity was gradually outgrowing them. The dramatic change in European attitudes toward dueling, slavery, and religious persecution over the last century led him to believe that "human nature is changing out of all recognition" and that man "is losing those psychological attributes that go with the employment of violence." He predicted that these changes would not take "thousands of years" to complete; it at least was not "impossible" that European attitudes toward war would evolve dramatically within "half a generation."

Angell's book was enormously influential. It drew widespread praise from British intellectuals, was translated into nearly a dozen languages, and spawned numerous university study groups devoted to discussing its ideas. Although later commentators would often characterize Angell as a starry-eyed idealist whose theories were contradicted by the First World War (e.g. Tuchman, 1962), his predictions were more long-term than this reputation suggests. He explicitly stated that "war is not impossible," that national defense remained a necessity, and that the naval arms race between Britain and Germany might tip over into war. His view, essentially, was that modern economic trends were gradually creating an *opportunity* for peace. To take advantage of this opportunity, it was still necessary to replace an outdated mindset about the value of war and introduce alternative methods of dispute resolution. A lasting peace was achievable, but remained to be achieved.

2.5 Interwar Thought

At least initially, the outbreak of the First World War did surprisingly little to suppress expressions of peace optimism. Many peace advocates became vocal supporters of the Allied cause, without abandoning their optimism. A common view in Britain and the United States was that the war was ultimately the result of an anachronistic spirit of militarism in Germany. If Germany could be defeated decisively, then there would be hope for a lasting peace in Europe. In a widely read sets of essays, the author and notable peace advocate H.G. Wells argued that this new war would be the "war to end all war" (Wells, 1914). Shortly after the entry of the United States into war, former president William Howard Taft served as the keynote speaker for a major national convention with the theme "Win the War for Permanent Peace" (League to Enforce Peace, 1918). When hostilities ultimately came to their end, Britain's Prime Minister, Lloyd George, would reaffirm this optimistic interpretation of the war: "I hope we may say that thus, this fateful morning, came to an end all wars."[1]

Despite its evident limitations and flaws, the post-war establishment of a League of Nations – a long-standing goal of the peace movement – offered another sustaining source of optimism. In the interwar period, there was a proliferation of schemes to expand and improve the League in order to ensure peace. Proposals to grant the League a global monopoly over airpower, which it could use to enforce peace through the threat of strategic bombing, became especially popular in Britain and France (Zaidi, 2011). These proposals were discussed extensively at a 1932 League of Nations conference on disarmament and received the support of prominent figures such as Winston Churchill. However, as Europe and Asia slid into the Second World War, the spirit of at least cautious optimism that motivated such proposals once again receded.

[1]It is perhaps worth commenting again, at this point, that rhetoric of peace optimists often suggested either an indifference to wars outside of Europe and North America or, especially in the case of colonial wars, a failure to conceptualize them as "wars." It is difficult to interpret the hopes surrounding the First World War otherwise.

E.H. Carr's *The Twenty Years' Crisis*, a foundational work of international relations theory published in 1939, memorably disparaged much interwar political thought as "utopian" (Carr, 1939). He instead argued in favor of a "realist" perspective on international politics that treats war as something closer to a timeless feature of the world, which can be managed but is unlikely to be expunged. Although Carr did not actually dismiss the possibility of a peaceful future, he considered it, at minimum, something that lay much further in the future.[2] As the academic field of international relations took shape, in the years following the Second World War, Carr's analysis would prove to be very influential (Schmidt, 2002).

2.6 Early Cold War Thought

The Cold War was a low point for peace optimism. The world now seemed to be growing increasingly violent. The twentieth century was commonly perceived as the most violent yet in human history. Furthermore, as the new bipolar world order took shape, belief in an inexorable trend toward interdependence and global democracy became very difficult to sustain.

The peace movement did not die out, and bold proposals for international institutions did not immediately disappear. For instance, there were fairly serious discussions between the United States and Soviet Union regarding the international control of nuclear weapons (Zaidi and Dafoe, 2021).

However, among peace advocates, desperation rather than optimism had become the dominant tone. It is difficult to find anyone arguing that pacifying forces would gradually expunge war from the world. The philosopher and longstanding peace activist Bertrand Russell's 1951 essay for *The Atlantic*, "The Future of Man," is one particularly extreme

[2]For instance, Carr closes out *The Twenty Years' Crisis* by musing that evolving economic institutions and moral progress could *eventually* lay a foundation for a more far-off peace. Reflecting on his own positive vision, he writes: "This, too, is a utopia. But it stands more directly in the line of recent advance than visions of a world federation or blue-prints of a more perfect League of Nations. Those elegant superstructures must wait until some progress has been made in digging the foundations."

illustration of this turn toward pessimism (Russell, 1951):

> *Before the end of the present century, unless something quite unforeseeable occurs, one of three possibilities will have been realized. These three are: (1) The end of human life, perhaps of all life on our planet. (2) A reversion to barbarism after a catastrophic diminution of the population of the globe. (3) A unification of the world under a single government, possessing a monopoly of all the major weapons of war. I do not pretend to know which of these will happen, or even which is the most likely...*
>
> *I think it is plain that the mutual suspicions between Russia and the West make it futile to hope, in any near future, for any genuine agreement [to establish a world government]. Any pretended universal authority to which both sides can agree, as things stand, is bound to be a sham, like UN. Consider the difficulties that have been encountered in the much more modest project of an international control over atomic energy, to which Russia will consent only if inspection is subject to the veto, and therefore a farce. I think we should admit that a world government will have to be imposed by force.*

This style of thought, while still reflecting a deep interest in the possibility of perpetual peace, could hardly be classified as an expression of "utopianism."

Due in part to Carr's influential attack on utopianism (later called "idealism"), the arguments of early twentieth century scholars like Norman Angell now also had very limited purchase in the new academic discipline of international relations. Whereas the work of earlier scholars was suffused with notions of evolution and progress, post-war scholars tended to discuss war in more timeless terms. While many scholars did advocate for a middle way between "realism" and "idealism" (e.g. Herz, 1951), anything approaching the notion of permanent peace was now distinctly outside the mainstream.

2.7 The Return of Peace Optimism

As the Cold War dragged on, then came to its surprisingly bloodless end, peace optimism once again began to emerge. Although it still was not an especially mainstream position, a number of notable works by political scientists raised the possibility that interstate war was on its way out. John Mueller's 1989 book *Retreat From Doomsday* argued that war between developed countries was becoming "obsolete" (Mueller, 1989). A highly cited article by James Lee Ray, published in the top international relations journal, suggested that, as in the case of slavery, moral progress may be causing war to disappear (Ray, 1989). Francis Fukayama's blockbuster book *The End of History* raised the possibility of a global peace founded on the spread of liberal democracy (Fukayama, 1992).

Three key factors help to explain the return of peace optimism at this time. First, after forty years of anxiety and dire predictions, many scholars believed that the failure of the Cold War to turn hot required an explanation (Gaddis, 1986). Early statistical research on trends in warfare also suggested the possibility of a general post-WWII decline in the rate of interstate war. Furthermore, historical research helped to correct a misperception that the twentieth century had overall been unprecedentedly violent. The notion that war might be in decline no longer appeared absurd.

Second, early statistical research had also begun to provide support for some of the central theories of older peace optimists. Joint democracy appeared to robustly predict lower rates of war (Maoz and Abdolali, 1989). By the end of the nineties, there was also substantial statistical evidence that trade and international institution membership predict peace (Oneal et al., 1996; Russett et al., 1998). An increasingly vast theoretical literature emerged to offer causal explanations for these links, combining a mixture of old arguments with new ones (such as the new argument that pairs of democracies suffer less from information asymmetries). The Kantian triangle was back. It became fairly common for researchers in this area to suggest, at least cautiously, that the world may be headed in a substantially more peaceful direction. As Eric Gartzke and Dominic Rohner put the point in one later paper, discussing

their regression results: "World peace is a long way off, but you may be able to see some of it from the results reported here" (Gartzke and Rohner, 2010).

Third, the collapse of the Soviet Union surely made it easier to imagine that these forces for peace would continue to spread across the globe. At least for a brief time, this momentous event may also have fostered a more general spirit of optimism about the future, a soft echo of the often utopian optimism that technological and social progress fostered in the nineteenth century.

To be sure, peace optimism was still far from a majority viewpoint. Few if any scholars were as bold as John Mueller. There were also a number of influential works, particularly John Mearsheimer's "Back to the Future," Robert Kaplan's "The Coming Anarchy", and Samuel Huntington's *The Clash of Civilizations*, arguing that the end of the Cold War would spawn new conflicts rather than facilitating peace (Mearsheimer, 1990; Kaplan, 1994; Huntington, 1996). Nonetheless, optimism about the future of war was no longer a fringe phenomenon.

2.8 The New Optimists

In 2011, the psychologist Steven Pinker published a bestselling book arguing that nearly every form of violence has been in decline for centuries (Pinker, 2011). The longest chapter in the book focuses on interstate war. Although the chapter is largely a work of synthesis, drawing particularly on work published by John Mueller and Jack Levy in the 1980s and 1990s, it caused a surge of interest in the claim that war has been gradually disappearing. More than any other scholar, Pinker has been responsible for the most recent rise in peace optimism.

To make his case, Pinker graphs the number of wars initiated in Europe between 1400 and 2000 and the number of wars involving great powers between 1500 and 2000. In both cases, to the naked eye, there appears to be a very noisy but nonetheless downward trend.

Pinker characterizes the world wars as random fluctuations, which temporarily obscured this long-run trend. He also lists a number of more qualitative or anecdotal observations that suggest a trend toward peace: the United States and Soviet Union successfully avoided war, the taboo against conquest appears to have grown over time, and developed states currently have especially peaceful relations with one another. Many pairs of states also now seem to take peace for granted, for instance by maintaining undefended borders and failing to seriously prepare for invasions by their neighbors, in a way they might not have in previous centuries. Citing statistical research on the correlates of war, Pinker argues that this apparent trend is largely explained by the Kantian triangle. In his view, however, the most fundamental causal factor has been a general change in attitudes toward violence. Like many previous writers, Pinker illustrates this change in attitudes by discussing the decline of dueling, slavery, torture, and other forms of violence beyond war itself.

Importantly, Pinker does not merely analyze the history of war. Echoing many earlier writers, he also raises the prospect of a more peaceful future, suggesting that interstate war may all but completely disappear in the long run:

> *[If this theory is correct] interstate war among developed countries would be going the way of customs such as slavery, serfdom, breaking on the wheel, disemboweling, bearbaiting, cat-burning, heretic-burning, witch-drowning, thief-hanging, public executions, the display of rotting corpses on gibbets, dueling, debtors' prisons, flogging, keelhauling, and other practices that passed from unexceptionable to controversial to immoral to unthinkable to not-thought-about during the Humanitarian Revolution.*

Since the publication of *Better Angels*, a number of other books and papers have presented variations on its core argument. For instance, Azar Gat, while agreeing with Pinker on many points, highlights a broader range of pacifying forces and suggests that many of them are downstream of industrialization (Gat, 2017). Gat ultimately arrives at a

similar position of "cautious optimism": although the risk of large-scale war is still all too real, and although we cannot be confident that war will continue to decline, we should also recognize that its present decline rests on an "extremely powerful trend" that can only be derailed by "extraordinary forces."[3]

Other notable and recent books arguing for at least a cautious form of peace optimism include John Mueller's *The Remnants of War* (Mueller, 2004), Joshua Goldstein's *Winning the War on War* (Goldstein, 2012), and the historian Ian Morris's *War! What is it Good For?* (Morris, 2014).

2.9 The New Pessimists

The arguments of Pinker and other "new optimists," although highly influential, have also attracted significant criticism (Mearsheimer, 2013; Cirillo and Taleb, 2016). Among their assorted critics, Bear Braumoeller has been by far the most thorough. His recent book *Only the Dead* – a reference to George Santayana's comment that "only the dead have seen the end of war" – is a sustained attempt to refute Pinker's claim that interstate war is in decline (Braumoeller, 2019). It is worth explaining this critique in some detail.

At its core, Braumoeller's critique is fairly simple. First, he notes that Pinker simply eyeballs graphs rather than performing statistical tests for the existence of a significant downward trend. This is problematic, he argues, because it is easy to see patterns in random noise. Since the number of wars per year is not normally distributed, as Pinker himself notes, it is especially difficult for the naked eye to distinguish between genuine and illusory trends.

Second, Braumoeller notes that Pinker does not adjust for changes in the number of states within his samples. A reduction in the number of European wars per year is less interesting, for example, if this decrease merely reflects a decrease in the number of European states. Braumoeller argues that it is necessary to divide the number of wars in each year by

[3]Gat (2017) is more relevant to this thesis than any other optimistic work, due its focus on industrialization; for pre-industrialized societies, industrialization is effectively synonymous with economic growth, and, for the most part, he could equally well have framed his arguments in the language of growth.

a measure of "politically relevant dyads."

Third, Braumoeller notes that Pinker considers declines within a limited Eurocentric sample.[4] His two primary datasets are a dataset of wars between European states and a dataset of wars involving (overwhelmingly European) great powers. Braumoeller argues that it is much more natural to analyze the Correlates of War dataset, which aims to include *every* interstate conflict since 1816.[5] If conflict has declined within Europe, but not declined overall, then this would seem to dampen the case for peace optimism. Such a finding might simply reflect a shift in alliance structures, with European states banding together in opposition to outside forces, rather than reflecting a meaningful trend toward global peace.

When Braumoeller plots the number of new interstate conflicts per year, adjusts for the number of politically relevant dyads, and applies a change-point algorithm to look for statistically significant rate changes, he finds that the global rate of conflict between 1816 and the present bounces up and down erratically. In the language of time series analysis, it exhibits *cycles*. If there is any underlying trend beneath these cycles, then this trend at least is not obvious. For example, the post-Cold War period is seemingly about as peaceful as the post-Napoleonic-War period that occurred more than a century earlier.

Braumoeller argues that changes in "international order" can have a large impact on the rate of war. For instance, the emergence of the "Concert of Europe" system may explain why nineteenth century Europe became so peaceful after the Napoleonic Wars. Similarly, the emergence of a unipolar world order after the end of the Cold War may have allowed the rate of war to fall. The rate of war bounces up and down, as new international orders emerge and dissolve. In his view, modern peace optimists tend to misinterpret recent one-off declines as evidence of a long-run trend that probably does not exist.

Beyond Braumoeller's revisionist analysis, the past several years of international poli-

[4]Eurocentrism has been one of the most common charges made against Pinker's work (Micale, 2018).

[5]In practice, as I will discuss in the next chapter of this thesis, the Correlates of War dataset is also initially Eurocentric. Its coding rules imply that the only two non-European states that existed in 1816 were the United States and Ottoman Empire. Effectively, it begins as a dataset of intra-European wars and only gradually expands to include wars beyond Europe's coasts. This issue means that it is actually fairly difficult to interpret trends in the aggregate rate of war.

tics have not been favorable to peace optimists. In the past year, large-scale war has returned to Eastern Europe. Other dramatic developments of the past decade include the outbreak of a war between Armenia and Azerbaijan and the outbreak of major civil wars in Syria and Yemen, which both attracted substantial participation by foreign powers. The major focus of US foreign policy continues to be great power competition with China; a war with China in the coming decades, triggered by an invasion of Taiwan, is seemingly a genuine possibility. The political moment no longer lends itself so well to expressions of optimism.

2.10 Conclusion

Expressions of peace optimism (the hope or expectation that interstate war will all but disappear) can be found as far back as the Axial Age. However, the modern variant of peace optimism largely took shape in the late eighteenth and early nineteenth centuries. Since this time, despite significant intellectual and empirical progress, the core elements of peace optimism have remained highly consistent. Some combination of democratization, rising trade, strengthened international institutions, moral progress, or (less often) one of several other factors is expected to gradually reduce the incidence of war. An apparent recent or long-run decline in the rate of war is interpreted as evidence that the key pacifying process is already underway. Social or technical progress in other domains – such as the abolition of slavery, the invention of the telegraph, or the gradual elimination of global poverty – is also often taken to support the general notion that progress on problems as seemingly intractable as war is plausible.

On one cynical interpretation, the repeated re-emergence of this package of ideas mainly reflects a failure to learn from history. Optimists draw excessively strong inferences from short-run trends and fail to pay enough attention to all the people who have mistakenly drawn the same inferences before them. This cynical interpretation is consistent, for example, with Braumoeller's recent analysis of long-patterns in interstate war.

However, this cynical interpretation is not the only possible one. On another interpretation, it is instead the repeated *disappearance* of peace optimism that reflects an excessive fixation on recent events and short-run trends. Consider, for instance, the case of Norman Angell. Angell explicitly indicated that future war with Germany was both plausible and consistent with his theories. Although economic interdependence was gradually creating greater incentives to avoid war, he argued, states had not yet adapted to these new incentives. Therefore, while the world wars certainly constituted evidence against Angell's theories, it was probably an overreaction for other intellectuals to nearly abandon them for decades. More obviously, given that democracies did not fight each other in either world war, these wars should have done relatively little to reduce optimism grounded in the spread of democracy. Communism was likewise too new and too untested for its initial rise, during this period, to justify abandoning the expectation that democracy would continue to spread in the longrun.

When interpreting historical trends and events, pessimists have not always paid enough attention to the question: "If the world were slowly headed toward perpetual peace, exactly how surprising would this observation be?" The answer will depend on both the particular causal mechanism that is supposed to produce a peaceful future and the nature of the stochastic process that produces random variation in the rate of war.

There is still a need, it seems to me, for more rigorous and systematic work assessing the possibility of a peaceful future.

First, moving beyond existing work, we should try to clarify exactly we should expect to see if there *was* a downward trend. Is it at all possible, for example, for a genuine longrun trend to be obscured by random motion and by data quality issues? To answer this question, we need to perform further statistical tests, run simulations, and, to some extent, clarify what exactly it means for there to be a "trend." The next chapter of this thesis will aim to carry out some of this additional work.

Second, we should try to clarify and systematically investigate the set of possible

causal pathways to a peaceful future. We should list them, taxonomize them, and assess their plausibility. We also should not take for granted that every plausible pathway has already been proposed. Then, once we have our list of candidates, we should ask: "If we are actually traveling along this pathway, what should we expect to see?" This work has not yet, for the most part, been done. The later papers will take a step in this direction by investigating one particular pathway that strikes me as underappreciated: the pathway leading from economic growth to peace.

We will never be able to predict the future of war with any great confidence. But we may still be able to distinguish the likely from the unlikely. Barring that, we can at least hope to distinguish the plausible from the absurd.

The papers that follow are attempts to nudge us, at least, a little closer toward this modest goal.

Chapter 3

Will the Future Be Peaceful? How Would We Know?

3.1 Abstract

This paper considers the question: "How likely is it that the future will be much more peaceful than the past?" I argue that the question merits serious study. Even though the field of International Relations will never converge on anything like a definitive probability estimate, it can still make progress by clarifying and compressing the *spectrum of reasonable disagreement* about the likelihood of a peaceful future. I outline a methodological approach that involves (1) forming an initial forecast purely on the basis of historical trends, then (2) modifying the forecast by considering possible causal explanations for higher or lower rates of conflict. I then produce an initial forecast, which suggests that a peaceful future is unlikely but not implausible. One source of hope is that a meaningful downward trend in the risk of conflict seems to have started in the mid-20th century, although we cannot easily predict how long this trend — if it is in fact real — will last. Modifying the initial forecast by incorporating causal considerations is left for later chapters and later work.

3.2 Introduction

3.2.1 The Peaceful Future Hypothesis

Will the future be much more peaceful than the past?

The field of International Relations emerged, in part, from a lineage of efforts to grapple with this question.[1] This lineage includes eighteenth century philosophers, such as Bentham (1789) and Kant (1795), who raised the possibility of a world without war; nineteenth century peace advocates, who campaigned for war's elimination; and early twentieth century theorists, most prominently Angell (1913), who developed more sophisticated arguments for believing a peaceful future was possible. Ultimately, the first person to be awarded the title of Professor of International Politics — in 1919 — was Alfred Eckhard Zimmern, a key figure in the design of the League of Nations and someone who had hoped the League would play a role in bringing an end to war (Markwell, 1986). For many years, the central motivating impulse for the young field was a desire to understand whether and how peace could be achieved.

Throughout this long history of optimism, the main arguments for the possibility of peace were relatively stable. Optimists typically invoked the pacifying effects of some combination of factors: economic entanglements, representative political institutions, international coordination and new international institutions, or intellectual and moral progress. Optimists were also often encouraged by the unusually low rate of war in nineteenth century Europe, by recent examples of dramatic social progress in other domains (such as the abolition of slavery), and by increasingly serious political interest in international governance (such as the creation of international law). The First World War of course dampened optimism, but did not destroy it. During the war, for example, idealistic rhetoric regularly framed the conflict as a "war to end war." There was still a notion that war might ultimately be on its way out, with the First World War serving as one great last gasp.

[1] See Chapter 2.

Although optimism continued to drain away during the interwar period, the possibility of a more peaceful future remained central to the field of international relations. The formative mid-century debates between so-called "realists" and "utopians" were, in large part, debates about whether a peaceful future was possible and what its preconditions might be.

Ultimately, though, the field largely moved on from debating the possibility of a peaceful future. The cataclysmic Second World War played an important role in this change. For many scholars, the question of whether peace was on its way out seemed to have been answered "no."[2] The field also, over time, gained a greater appreciation of the difficulty of forecasting. Whereas early scholars often made predictive claims, later scholars increasingly restricted themselves to descriptive and causal claims. Efforts to predict the future came to be regarded as somewhat unscholarly.

However, the field's interest in the *Peaceful Future Hypothesis* – the hypothesis that the future will be far more peaceful than the past – has never fully disappeared. Scholars studying historical trends in war or the causes of war will sometimes give themselves leeway to comment on the hypothesis. Even in these cases, though, the comments are typically confined to the concluding sections of the papers or books where they appear.[3] The hypothesis is seldom the *focus* of a major work. There sometimes seems to be a sense that staring too directly at it is a little like staring too directly at the sun.

This paper will take an unusually direct look at the Peaceful Future Hypothesis. In examining it, I will have three main aims: *First*, I will try to clarify the hypothesis and show that it is worthy of serious study. *Second*, I will try to create a helpful "map" of the hypothesis: this means clarifying what we already know, thanks to existing scholarship; how these findings relate to each other and bear on the hypothesis; and what work remains to be done. *Third*,

[2]This shift may have flowed substantially through generational turn-over, though, rather than older optimists about peace fully abandoning their optimism. For example, Zimmern's final book was published in 1953 and had the idealistic title *The American Road to World Peace* (Zimmern, 1953).

[3]For example, both Pinker (2011) and Braumoeller (2019) discuss historical trends in war across dozens of pages, but discuss the forward-looking implications of their findings across only a handful of paragraphs.

I will ultimately try to make some progress in assessing the hypothesis: my hope is that the reader will leave with an improved sense of how plausible or implausible a peaceful future actually is.

Over the course of the paper, to further these aims, I will produce an *intermediate forecast* of the likelihood of a peaceful future. This intermediate forecast will be based purely the analysis of historical trends in the rate of conflict. Further analysis, which incorporates considerations about the causes of war, would be needed to arrive at a more holistic overall forecast. Still, the intermediate forecast on its own suggests that a peaceful future is an unlikely but not implausible possibility. The basic positive case for a peaceful future is: a noisy trend toward peace seems to have started in the mid-twentieth century, and, since it has lasted for a long time so far, we at least should not be too surprised if it carries on for a long while longer.

The intermediate probability estimate I will ultimately arrive at is 20%. This particular number should not be taken very seriously: since it rest on many contestable judgment calls, meaningfully higher or lower numbers would also be defensible. Even so, in my view, it would be difficult to justify an intermediate forecast that is *dramatically* more or less confident.[4]

3.2.2 Overview of Paper

This paper is divided into six sections.

The first section attempts to clarify the Peaceful Future Hypothesis. It offers two operationalizations of the hypothesis that are meant to enable more concrete forecasts and debates. The primary operationalization of the question is: "Will the average rate of conflict between 2050 and 2300 will be at least 90% lower than it was between 1850 and 2000?"

The second section offers a general defense of the value of forecasting and future-oriented work in political science. It outlines the principles of probabilistic forecasting and

[4]For example, I believe it would be difficult to defend a number as low as 1% or as high as 60%.

summarizes key practical lessons from empirical studies of successful forecasts. It also clarifies what we should expect intellectual progress in assessing the Peaceful Future Hypothesis to look like. We should not expect the field to ever converge on anything like a definitive probability estimate. However, we can still hope to clarify and narrow the *spectrum of reasonable disagreement* about the likelihood of a peaceful future.

The third section outlines a two-step process for estimating the likelihood of a peaceful future. The first step involves generating an "outside view" forecast based on analysis of long-run trends in the rate of conflict. The second step involves modifying this initial forecast by evaluating different causal mechanisms that could lead the rate of conflict to overshoot or undershoot the current trend — or that could give us reason to expect the current trend to continue.

The fourth section walks the reader through the process of generating the initial outside view forecast, noting the key judgment calls that need to be made. This section draws especially heavily on seminal work by Braumoeller (2019). In line with Braumoeller's work, it focuses specifically on the rate of "interstate militarized disputes involving the reciprocated use of force" as the key variable to forecast. This section's main innovation is the use of a modified "geometric Brownian motion" model to search for a trend in the historical rate of conflict. A key finding is that the data we have is consistent with a wide range of possible underlying trends, including a substantial rising trend or a substantial declining trend. Furthermore, due to inherent randomness, each possible underlying trend is consistent with a wide range of possible future trajectories. Even accounting for these sources of uncertainty, though, our model assigns a very small probability to a peaceful future.

The fifth section then adjusts the initial forecast to account for three more fundamental sources of uncertainty: the possibility of trend breaks (in the past and in the future), the possibility that our dataset is flawed, and the possibility that our model does not accurately describe how the rate of conflict evolves. Accounting for possible trend breaks turns out to be especially consequential. We find strong evidence of a downward trend beginning in

the mid-twentieth century; when extrapolated forward, this more recent trend implies that a peaceful future is actually very likely. However, once we account for the possibility of further trend breaks in the future, our overall probability estimate falls back down substantially. Accounting for data uncertainty and model uncertainty leads to smaller changes in our probability estimate. We ultimately arrive at an **overall outside view forecast of 20%**. However, I also stress that — to produce this number — we have needed to make many judgment calls. Meaningfully higher or lower estimates could also be considered reasonable.

The sixth section briefly outlines the path forward, to move away from our purely "outside view" forecast and incorporate causal considerations. However, I leave this additional analysis for later chapters and future work.

3.3 Clarifying the Peaceful Future Hypothesis

This paper is preoccupied with the question of whether the future will be far more peaceful than the past. In slightly more specific terms, the relevant hypothesis is:

> **Peaceful Future Hypothesis:** There will be (or has already started to be) an extremely large and long-lasting decline in the rate of interstate conflict, relative to its historical level.

Over the past couple hundred years, as summarized in the previous chapter, many authors have expressed views on this hypothesis. Of course, though, the hypothesis is normally described in fairly imprecise terms.

One informal approach to describing a peaceful future, which has been employed by a number of optimists about peace, is to draw an analogy with the decline of slavery (Ray, 1989; Mueller, 2004). On a broad definition of "slave" -- which includes indentured servitude, serfdom, and other forms of unfree labor -- the majority of the world's population may have been enslaved as recently as the eighteenth century (Eltis and Engerman, 2011). While

slavery still exists, there is also no doubt that its global prevalence has fallen dramatically.[5] Furthermore, even if there are important resurgences, it is now nearly inconceivable that – absent other radical changes in the world – the rate of enslavement will ever return to anything close to its pre-industrial peak. One informal description of a peaceful future, then, is one in which interstate war eventually "goes the way of slavery."[6]

For the sake of reducing ambiguity, though, it is also useful to have a more specific version of the question to discuss. In this paper, the main formulation I will focus on is:

> **Peaceful Future Hypothesis (Average Rate Version):** The *average* rate of interstate conflict between 2050 and 2300 will be at least 90% lower than the *average* rate of conflict between 1850 and 2000.

I will also sometimes consider another secondary formulation:

> **Peaceful Future Hypothesis (Peak Rate Version):** The *peak* rate of interstate conflict between 2050 and 2300 will be at least 90% lower than the *peak* rate of conflict between 1850 and 2000.

These high (90%) thresholds are meant to ensure that our discussion stays focused on a genuinely radical proposition, which is in the spirit of historical discussions of the possibility of a peaceful future. At the same time, these thresholds also allow a meaningful level of conflict to persist in the world.

The choice to focus on such a long future period has a couple of motivations. One motivation is to make sure that major changes in rates cannot simply reflect random fluctuations. Another motivation is that, if we want to give a fair hearing to arguments for peace,

[5] By the end of the twentieth century, the rate of enslavement had likely fallen far below 1% (Bales et al., 2009).

[6] This analogy to the abolition of slavery dates back to the nineteenth century. The successful British abolitionist movement, which was intertwined with the peace movement, offered inspiration and a model for those hoping to eliminate war (Cortright, 2008). It has often been noted that, like war, slavery is another practice that dates back to at least to the dawn of civilization and was once regarded as inevitable (Ray, 1989).

it seems that we should give any existing trend toward peace enough time to run its course. At the same time, choosing longer periods or periods that lie further in the future will tend to make the question both *more difficult to analyze* and *less practically relevant*.

Overall, these operationalizations are of course fairly arbitrary. Nonetheless, settling on specific operationalizations will tend to make our analysis of the hypothesis clearer and more productive. We will actually continue to make the operationalizations more specific, over the course of the paper, as some of its ambiguities (such as how the "rate of conflict" should be calculated) become more salient.

For some of the aims of this paper, the operationalizations we choose actually will not matter very much. The paper is concerned, in insignificant part, with methodological matters. These fundamental methodological issues involved in assessing the Peaceful Future Hypothesis do not change much if we use one operationalization or another.

3.4 Forecasting and Political Science

When scholars express views on the Peaceful Future Hypothesis, they are attempting to predict the future. In other words, they are making *forecasts*.

There is a modest but growing International Relations literature on forecasting conflict (e.g. Choucri, 1974; Ray and Russett, 1996; Schneider et al., 2011; Hegre et al., 2017). For instance, in one particularly innovative paper, Hegre et al. (2013) have produced a set of forecasts of how likely individual countries are to experience civil conflicts between 2010 and 2050.

Nonetheless, political scientists still often regard forecasts with suspicion – and not without reason. Rigorous studies suggest that even domain experts are habitually overconfident when forecasting future political developments (Tetlock, 2017). The inherent complexity and opacity of political phenomena also place important limits on the level of foresight it is possible to achieve. For example, it is hard to see how any rational analyst in 1945 could

have *confidently* predicted that nuclear weapons would not be used again within the century.

Nonetheless, political scientists should not regard forecasting questions as unworthy of their attention. In this section, I first present a brief defense of future-oriented political science research. I then summarize empirical literature on best practices for forecasting. This literature suggests a rough path forward for scholars grappling with the Peaceful Future Hypothesis.

3.4.1 In Defense of Forecasting

Forecasts of political trends and events are often overconfident. In addition, in many cases, it simply is not possible to make rationally confident predictions – no matter how much effort you put in. We can study for as long as we want, but even years of analysis will not let us confidently predict which side of a coin will face up after it's flipped. Likewise, for many political questions, there is no achievable level of expertise or analysis that will be sufficient to justify confident predictions.

However, these observations do not imply that political scientists should simply ignore forecasts. In fact, we have good reason to devote effort toward informing, critiquing, or producing forecasts.

First, the answers to many political forecasting questions simply *matter* a great deal. For example, the likelihood that the world will avoid nuclear war in the next few decades has significant practical consequences. A 99% chance of avoiding nuclear war would have quite different implications than a 50% chance. The latter probability would suggest, as one grim but straightforward implication, that young people should save less for retirement. It would also suggest that governments should be more willing to deprioritize foreign policy objectives - perhaps including some economic objectives or even human rights objectives - that are in tension with great-power peace. Philanthropists would also have good reason to increase their funding of nuclear security initiatives. Different forecasts have meaningfully different implications.

Second, regardless of whether the discipline engages with forecasting questions, people will inevitably continue to form and act on beliefs about them. In general, when people make decisions, they must also make predictions. Someone deciding whether to take an umbrella is making at least a fuzzy, implicit prediction about the likelihood of rain. In the same way, when European policymakers were deciding military funding levels and energy policy decisions – in the years before the Russo-Ukrainian War – they were making at least fuzzy, implicit predictions about the likelihood of large-scale war returning to Europe. As Philip Tetlock puts it: "All policies are predictions."[7] Therefore, political scientists are not in a position to prevent political forecasts from being made. The only real question is whether or not the discipline should attempt to help people arrive at better answers.

Third, even if it is often unreasonable to make confident predictions, this does not imply that we cannot make progress on forecasting questions. To explain this point, it is worth keeping in mind that a set of forecasts can be judged in terms of both *confidence* and *calibration*.[8] A set of forecasts is said to be "well-calibrated" if predictions that are expressed with 50% confidence come true about 50% of the time, predictions that are expressed with about 60% confidence come true about 60% of the time, and so on. (See Figure 3.1.) Someone who is well-calibrated is, essentially, *appropriately confident* in the forecasts they make.

We can always attempt to achieve an appropriate level of confidence, even if – due to limited information or the world's inherent complexity – that appropriate level of confidence turns out to be quite low. Establishing that people should only be about 50% confident in some prediction, when they might otherwise have been overconfident, is actually a useful form of progress. *Calibration is a worthy goal, even when rational confidence remains out of reach.*

[7] Of course, even people who do not need to make decisions often take a strong interest in questions about the future. Popular articles and books about the future are ubiquitous.

[8] This is also known as "discrimination" or "sharpness."

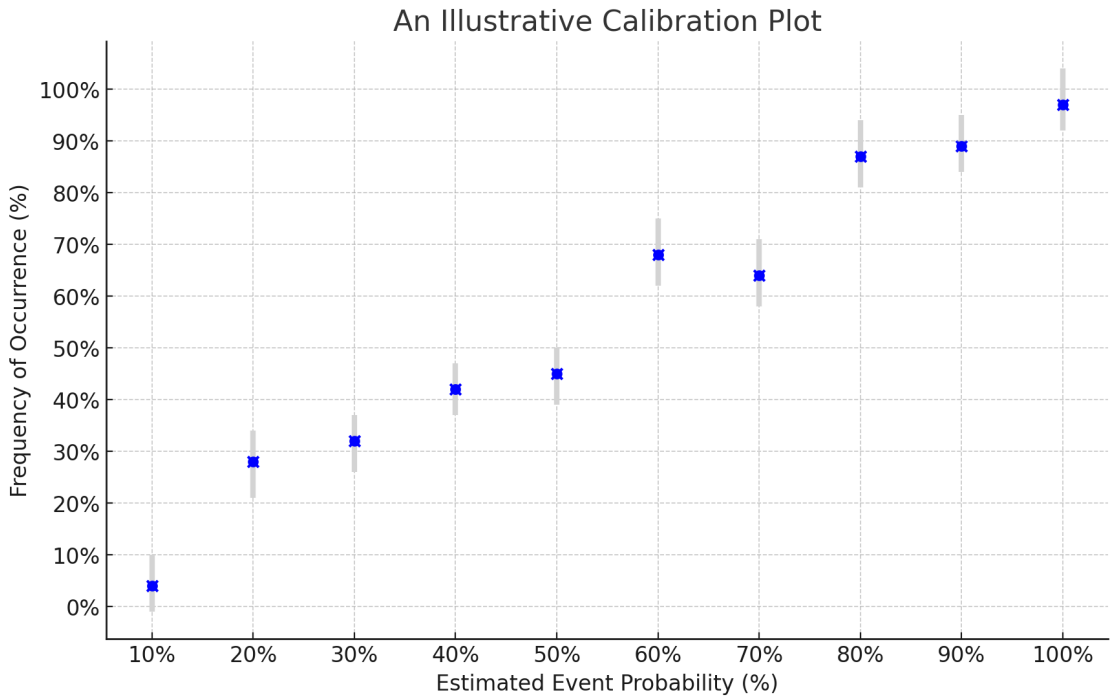

Figure 3.1: This plot illustrates a "well-calibrated" set of forecasts. For example, in this set of forecasts, events that are assigned a 30% probability ultimately occur approximately 30% of the time. If some person's or some statistical tool's past forecasts look like this when plotted, then this would count as evidence that they tend to make appropriately confident forecasts.

3.4.2 Forecasting: What Works

Empirical research confirms that it is possible for people to generate well-calibrated predictions, in various political and social domains, while at least sometimes expressing high levels of confidence. This research also offers important lessons about forecasting methodology, which are based on observations of strong and weak forecasters.

One key lesson is that strong forecasters often operate within the same rough methodological framework, which I will call the *Tetlock Framework*. We should expect the Tetlock Framework to provide a useful structure for attempts to grapple with the Peaceful Future Hypothesis.

I will now summarize the relevant literature and briefly describe the relevant framework.

3.4.2.1 Empirical Studies of Forecasting

Forecasting performance can be measured in a variety of ways. One commonly used metric, for example, is known as a *Brier score*. If someone makes a set of probabilistic predictions, then we can calculate their Brier score by comparing their predictions against observed outcomes.

A person's Brier score depends on how well-calibrated their predictions are and, secondarily, on how confident their predictions are. Confidence is factored into Brier scores for the simple reason that well-calibrated high-confidence predictions are more *useful* than well-calibrated low-confidence predictions. For example, a set of 99% confident predictions that come true 99% of the time will be more useful than a set of 50% confident predictions that come true 50% of the time. (If confidence was not factored in, then someone could conceivably earn a high score simply by assigning a 50% probability to every hypothesis they are asked about; this would not be very helpful.)

By using metrics such as Brier scores, researchers have empirically studied human forecasting ability in domains such as world politics. One influential approach is to have people compete in forecasting tournaments. These tournaments ask participants to make a large number of probabilistic predictions about what will happen within some window of time (e.g. the next year). Once the window has passed, researchers then calculate how well the various participants performed and attempt to identify the causes of high performance.

One reassuring finding from this research is that, in the domain of world politics, good forecasting is fairly possible (Tetlock, 2016). Some forecasters are able to maintain very high levels of calibration, of the type illustrated in Figure 3.1, while also at least sometimes putting forward confident predictions.[9]

[9]Of course, inevitably, most of our evidence concerns relatively short-run forecasts. We should expect, all else equal, that the difficulty of making confident predictions will tend to be higher the further the relevant event lies in the future. However, it would be surprising if the principles of good forecasting differed entirely between short-run and long-run forecasts. If there are techniques that allow forecasters to get a purchase on events that lie a few years in the future, we should seemingly expect these techniques to allow them to get at least some purchase on events that lie further in the future.

3.4.2.2 Inside View and Outside View Methods

Through their empirical studies, researchers such as Philip Tetlock have made progress in clarifying *how* some forecasters achieve high levels of performance. Some of Tetlock's most notable findings concern the ways in which strong forecasters combine two different kinds of forecasting methods, known as "inside view" and "outside view" methods. It is worth taking a moment to explain the nature of these methods.

Roughly speaking, *inside view* methods involve causal reasoning and are sensitive to the details of the case under consideration. For example, consider a commentator who is trying to predict how fast China's growth will be over the next decade. The commentator puts forward the following argument:

> *China's growth will probably decline a great deal, since Xi Jinping's desire to consolidate power will probably lead him to crack down increasingly hard on companies he perceives as being too powerful. These crackdowns will reduce private-sector innovation, which in turn will limit growth.*

This commentator's forecast is implicitly evoking a set of causal theories – linking the desire for authoritarian power consolidation to corporate crackdowns, linking corporate crackdowns to private-sector innovation, and linking private-sector innovation to growth – and applying them to the specific case of China. In this sense, the commentator is taking an "inside view" approach to forecasting.

Roughly speaking, in contrast, *outside view* forecasting methods do not rely on causal reasoning and are fairly insensitive to the details of the case under consideration. For example, consider an alternative commentator who puts forward this argument:

> *China's growth in the next decade will probably be fairly similar to what it was in the last decade, since* – usually – *a country's growth rate does not change much from one decade to the next. However, because there has been*

a downward trend in China's growth rate for a while, it will also probably end up being at least somewhat slower in the next decade.

This commentator's forecast is not really evoking any causal theories about economic growth. It is also largely ignoring the specific details of China's economic and political situation. In this sense, the commentator is taking an "outside view" approach to forecasting.

Outside view forecasting mostly relies on *reference class comparison* and *trend extrapolation*. Reference class comparison involves situating the relevant case within a larger set of similar cases, then asking what is *usually* true for cases within the set. Trend extrapolation involves simply noticing trends and extrapolating them forward.[10] In the above example, the second commentator engages in both reference class comparison and trend extrapolation.

3.4.2.3 The Tetlock Framework

In examining how strong forecasters combine inside view and outside view methods, Tetlock has noticed a significant two-part pattern. We can call this pattern the *Tetlock Framework*.

First, when possible, strong forecasters tend to rely on outside view methods as the foundations for their forecasts. Even very basic trend extrapolation and reference class comparison approaches can often outperform inside view approaches. For this reason, strong forecasters often *start out* by forming an initial prediction based on outside view methods. They will then *adjust away* from their initial "outside view prediction" if they encounter sufficiently compelling causal considerations.

Second, when strong forecasters do apply inside view methods, they often consider a mixture of causal theories rather than assuming a particular causal theory. For example, a strong forecaster attempting to predict China's growth might consider a range of possible

[10] In some sense, trend extrapolation is actually a particular variety of reference class comparison — where the reference class being looked at is a set of previous periods of time. Someone extrapolating a downward trend is forecasting that the variable of interest will probably decline in upcoming periods, because it has tended to decline in some set of previous periods.

causal factors – including China's political leadership, economic policies, and demographics – while also keeping in mind that there is a great deal of uncertainty about the causal relationships between these factors and economic growth.

Of course, there are also an enormous number of degrees of freedom involved in applying this general pattern. Ultimately, good forecasting inevitably involves a great deal of subjectivity, art, and sensitivity to methodological nuances.

However, if we want to tackle a difficult forecasting challenge – such as forecasting future rates of conflict – then the Tetlock Framework is one of the most useful sources of guidance we have. Strong forecasters tend to operate within the bounds of this broad framework.

3.4.3 What Intellectual Progress Looks Like

Of course, for most forecasting questions, it is in practice impossible to produce anything like a *definitive* forecast. For essentially any forecasting question, there will still be at least some room for reasonable people to disagree. For instance, no matter how analysis and debate we do, we will not be able to identify a single objectively rational level of confidence in the hypothesis that a Democratic candidate will win the next US presidential election.

However, when it comes to forecasts — or, in general, any topic where there is the potential for disagreement among experts — different possible views clearly lie along a *spectrum of reasonable disagreement*. It would clearly be *reasonable*, as an obvious example, to assign a 50% chance to a standard coin landing heads when flips. It would clearly be unreasonable, on the other hand, to assign a 40% chance to it landing heads, in the absence of any special knowledge about the coin. In the same way, assigning a 50% chance to a Democratic candidate winning the next election is at least ambiguously reasonable, but, at the time of writing, assigning a 99% chance to this possibility clearly would be unreasonably confident.

What intellectual progress tends to look like, then, is the spectrum of reasonable disagreement becoming narrower. In the case of the Peaceful Future Question, in particular, the field of international relations cannot hope to converge on a definitive likelihood estimate. The field can hope, however, to clarify and compress the spectrum of reasonable views.

3.5 How Can We Forecast the Future of Conflict?

3.5.1 A Two Stage Process

The Tetlock Framework suggests that – if we want to make progress in assessing the Peaceful Future Hypothesis – then we can break the forecasting problem down into something like a two-stage process.

First, we can form an initial "outside view" forecast by considering trends in conflict. These trends allow us to make predictions, by extrapolating them into the future. If the risk of war has already been declining for a long time, for example, then perhaps this decline will also continue for a long time.

Second, we can then adjust our initial forecast by considering a range of "causal stories" about the future of war. For example, many scholars believe that when states become democracies they become less likely to fight each other. Some scholars also believe that democracy will become more prevalent over the next century. These two assumptions suggest a *causal story* about the decline of war: perhaps war will decline because democracy will spread and make states less likely to fight each other. We can evaluate this causal story by assessing both the core causal claim (that democracy causes peace) and the complementary forecasting claim (that democracy is likely to become more prevalent). If we ultimately find the story sufficiently compelling, then we may want to adjust the probability that we assign to a peaceful future away from what pure trend extrapolation would suggest.[11]

[11] Causal stories can also inform the probabilities we assign to various possible trends *already being in place*. The reason for this is that, in practice, it is often difficult to clearly identify historical trends from empirical data. For example, we will find that it is surprisingly difficult to say whether or not there is currently a

3.5.2 Plan for the Remainder of the Paper

For the remainder of the paper, I will walk through the first stage of this process. I will mostly leave the second stage for future work.

In walking through the first stage of the process, I will arrive at an intermediate probability estimate based on "outside view" analysis. Generating this estimate will require making many judgment calls. These judgments will be to a large extent *subjective* (since reasonable people could make different calls) but not *arbitrary* (since some calls would also clearly be less reasonable than others). The estimate I arrive at will therefore be far from authoritative, but should still serve as at least an example of a *reasonable* estimate.

Ultimately, as noted above, this exercise has three main goals. They are, in order of priority:

1. **Demonstrate tractability:** The first goal is to convince the reader that it is actually possible to make progress on the Peaceful Future Question (i.e. do work that clarifies and narrows the range of probability estimates that can be considered "reasonable").

2. **Map the forecasting problem:** The second goal is to provide a kind of "map" of the forecasting problem. This means clarifying what we already know, thanks to existing scholarship; how these existing findings relate to each other and to the overall forecasting problem; and what work still remains to be done.

3. **Clarify how likely a peaceful future is:** The third goal is to provide the reader with at least a rough and tentative sense of what kinds of estimates, for the probability of a peaceful future, can be considered reasonable.

downward trend in the rate of war.

3.6 Trend Extrapolation: Forming an Initial "Outside View" Forecast

3.6.1 Basic Approach

At first glance, trend extrapolation appears to be a fairly straightforward approach to forecasting the future of conflict.

There are three seemingly simple steps: First, we must pick a specific measure of the rate of conflict. (This will also allow us to make our forecasting question more specific.) Second, we must gather data on the value of that variable over time. Third, we must look for a trend in this variable and note what would happen if the trend continued.

As we will see, though, each of these steps is quite complicated in practice. Since Braumoeller (2019) provides the most thorough and influential analysis of trends in conflict, my discussion will frequently cite his work. I will align with his approach to the first two steps, while also engaging in a more extensive analysis of alternative options. I will then go substantially beyond his work in the third step.

3.6.2 Choosing a Metric for the Rate of Conflict

If we want to answer the question "Is there a downward trend in the rate of conflict?" – and, ultimately, "Will the rate of conflict decline?" – then we need to clarify what we mean by the rate of conflict.

In practice, this problem has two parts. First, we need to clarify what we mean by "conflict." Then we need to clarify what we mean by "rate." Both of these choices will turn out to be surprisingly complex.

3.6.2.1 Clarifying "Conflict"

3.6.2.1.1 Choosing a Severity Threshold: The Relevance/Practicality Trade-Off We need to decide on a severity threshold for classifying an event as a "conflict."[12] For example, we need to decide whether we are focusing only on the largest wars or whether we are focusing only on the largest wars or whether we are also including small-scale disputes.

The core problem facing us — and anyone who needs to choose a metric to forecast — is that we want to satisfy two criteria that sometimes conflict. These criteria are:

1. **Relevance:** What metric has the most practical significance?
2. **Practicality:** What metric (or proxy metric) lends itself best to trend extrapolation?

In our case, it is much more practical to analyze trends in very inclusive categories of conflicts (e.g. interstate disputes of any sort) than it is to analyze trends in very exclusive categories of conflicts (e.g. wars between great powers or wars that kill at least one-thousandth of the globe). Conflicts in these more exclusive categories are sufficiently rare that we would be working on datasets that essentially consist of a handful of non-zero data points scattered over centuries. This simply does not give us much to work with, either in identifying trends or — later on — in empirically analyzing causal relationships.

At the same time, we also care more about these more exclusive categories. Suppose, for example, that the rate of "interstate disputes" will double over the next century — but that major wars will essentially disappear. The disappearance of major wars would clearly be the trend that matters more. The trend in "interstate disputes" would be relatively unimportant on its own, if all of these disputes turned out to consist of (e.g.) fishing boat seizures or shots fired by individual disobedient soldiers.

Arguably, the specific probability that matters most is the probability of large-scale nuclear war. If a nuclear war were to occur, then billions of people could die. The probability

[12]The term "severity" sometimes has a technical meaning in international relations, referring to the total number of battle deaths resulting from a conflict. However, I do have in mind quite this specific of a concept.

matters. Of course, though, there is no meaningful empirical "trend" in the rate of nuclear war. The relevant dataset is a spreadsheet that starts in 1949 (the first point where multiple states possessed nuclear weapons) and then contains a zero for every year since then.

The usefulness of studying trends in more inclusive categories will depend on the extent to which the following assumption holds:

> **Constant Relative Probabilities Assumption:** The relative probabilities of conflicts of different severity levels are *roughly* constant over time. For example: if low-level disputes become twice as likely, then major wars become *roughly* twice as likely as well.

If the assumption holds, then trends in inclusive conflict categories are directly informative of trends in more exclusive categories. A downward trend in the probability of "interstate disputes" in general would imply a parallel downward trend in (e.g.) the probability of nuclear war, even if we cannot observe the latter trend.

Unfortunately, empirically, it is also hard to know if the Constant Relative Probabilities Assumption holds. The fact that it is hard to assess trends in the *absolute* probabilities of rare conflicts also means that it is hard to assess trends in their *relative* probabilities.

Ultimately, even if we are most interested in high-severity conflicts, the best approach may still be to *start out* by studying trends in a relatively inclusive category of conflicts. We can start out with the tentative assumption that the Constant Relative Probabilities assumption holds, then — once we have our initial forecast — probe this assumption more closely. We can look for signs that the Constant Relative Probabilities assumption is false and consider evidence and arguments that bear specifically on higher severity conflicts.[13]

[13]The most statistically sophisticated approach we can take to probing the Constant Relative Probabilities assumption is to: define a metric over conflict severity, model it as following a particular statistical distribution, and then attempt to detect changes in this distribution. Braumoeller (2019) takes this approach, using the portion of combatants killed in an interstate conflict as his severity metric and a power law as his assumed probability distribution. He finds no evidence of a change in the parameters of the distribution within the past two centuries. However, it would be difficult to detect a meaningful change even if there was one. The error bars on the parameter values are very large and, in practical terms, the test is mainly a test for whether the Constant Relative Probabilities assumption holds between low-intensity and medium-intensity conflicts.

In this paper, I will not carry out the second step. (Although Appendix A of this thesis will briefly comment on the path forward.) Instead, I will just focus on forecasting a *relatively* inclusive category of conflict. Following Braumoeller (2019), **I will focus on "militarized interstate disputes involving the reciprocated use of force."** These are disputes where one state uses force against another state and then that state uses force in return. Braumoeller (2019) argues that they represent a relatively meaningful step up the ladder of escalation, suggesting some level of willingness to move to a higher level of conflict if need be. (This is, essentially, an argument for suspecting that the Constant Relative Probabilities assumption might hold for them.) At the same time, they are already quite rare, and using a higher severity threshold would substantially increase the difficulty of detecting trends or studying causal relationships.

3.6.2.1.2 Including or Excluding Intrastate Conflicts So far, our discussion has been framed around interstate conflicts. However, this choice is not entirely obvious. Both varieties of violent conflicts are, of course, deeply important. It also does not necessarily matter much, for people who may be affected by violence, which side of a border the violence would flow from.

However, since the dynamics and drivers of these two types of war seem to be somewhat different, the trends for them can diverge. For example, in the second half of the twentieth century, the number of civil conflicts appeared to drift sharply upward at the same time that the number of interstate conflicts seemed to drift downward. If that is right, then it will be useful to separate the two varieties of conflict out.[14]

Furthermore, it may be easier to forecast interstate conflicts. The datasets that are

There are too few extremely high-intensity conflicts for changes in their underlying probabilities to easily affect the results.

[14]This separation approach does become more questionable the further back we look back in time and the further we move outside of Europe, since the line between the two varieties of conflict tends to become blurrier. The modern European-influenced state system and its conception of state sovereignty makes it unusually easy to place conflicts into one category or another. Nonetheless, even in the most distant past, categorization is still ultimately possible. For example, we can easily say that Rome's wars with Carthage were interstate wars and that Julius Caesar's war with Pompey was a civil war.

available for interstate conflicts tend to extend further back and contain a broader set of conflicts than the datasets for intrastate conflicts.

This is, essentially, a practicality argument for focusing — at least initially — on interstate conflicts. **For the remainder of this paper, I will continue to focus just on interstate conflict.** However, the broad approach taken in this paper can be applied to intrastate conflicts as well.

3.6.2.1.3 Disaggregating Conflicts Finally, we face a problem of disaggregating conflicts It is not entirely clear what should be counted as a "single conflict"

Part of the challenge is deciding on criteria for identifying start and end dates. The more important challenge, though, is deciding how to count entangled conflicts. For instance, given its complex global web of participants, should we count the Second World War as a single conflict or as many?

The approach this paper will take is to treat each pairwise conflict – each case where one state fights another – as its own conflict. This approach to counting, which Braumoeller (2019) also follows, corresponds more closely with what we care about. If we want to know whether the world is becoming more peaceful, then the Second World War *should* count for more than the Falklands War.

From a practical perspective, this approach to disaggregating conflicts also reduces the need to engage in arbitrary and subjective judgements about which dyadic conflicts should be grouped together. We do not need to worry, for example, about whether the Pacific and European components of the Second World War were truly part of the "same" conflict.

3.6.2.2 Clarifying "Rate"

Even once we have clarified what we mean by "conflict," we then need to clarify what we mean by "rate." There are three main approaches to defining the rate of conflict in a given year: (i) defining it as the number of conflicts, (ii) defining it as the number of conflicts per

state, and (iii) defining it as the number of conflicts per (politically relevant) pair of states.

3.6.2.2.1 Option 1: Number of Conflicts The simplest option is to define the rate of conflict as the number of conflicts that occur in a year.

However, this does not seem to be the best option. The problem here is that the number of states has changed a great deal over the past few centuries. Suppose, then, that we find that there has been a decline in the number of instate conflicts – but also a corresponding decline in the number of states. This finding would not necessarily imply that the world has been becoming "more peaceful" in a sense we care about.

3.6.2.2.2 Option 2: Conflicts Per State The next option is to divide the number of conflicts by the number of states in existence.

This option is attractive, because the metric corresponds to something intuitively important: the average number of conflicts that each state fights. A downward trend implies that someone living in the average country will, for example, experience fewer interstate conflicts in their lifetime. A world where the average country is more frequently at peace is, in an intuitive sense, a more peaceful world.

However, from a practical perspective, this metric may not be the best one to use when extrapolating the rate of interstate conflict. The problem here is that changes in the number of states also likely affect the number of conflicts-per-state. When the number of states grows, the number of *pairs* of states ("dyads") grows exponentially.[15] If we assume conflict probabilities are at least somewhat independent – such that each dyad's decision to fight has only a limited relationship with any other dyad's decision to fight – then the expected number of conflicts-per-states grows as the number of states grows. Therefore, trends in the number of conflicts-per-state may substantially reflect trends in the number of states.

[15]For example, consider an idealized world with just two states. In this world, there is only a single dyad that can fight. Now, imagine introducing one more state. Introducing this new state triples the number of dyads to three.

The number of states is - seemingly - unlikely to change much going forward. (Colonialism and decolonization were major drivers of changes in the number of states over the past few centuries.) This implies that, when we are doing trend extrapolation, we should probably try to "screen off" the portion of the trend that was driven by changes in the number of states. We should try to understand what the trend *would have been* if the number of states had not changed.

3.6.2.2.3 Option 3: Conflicts Per (Politically Relevant) Dyad A superior approach – for doing trend extrapolation – may therefore be to divide the number of interstate conflicts by the number of dyads.

However, unsurprisingly, there are also problems with this approach. One problem is that dyads, due to geography, are not easily able to fight each other or simply have no plausible reason to fight. For example, it is hard to imagine Vanuatu and Niger fighting a war. We should not expect the existence of such *politically irrelevant dyads* – such as Vanuatu-Niger – to substantially influence the number of conflicts that happen in a given year.

What we may want to do, then, is to divide the number of conflicts by the number of "politically relevant dyads." This is the approach that Braumoeller (2019) takes. To determine the denominator, he assigns each dyad a political relevance score. He calculates these scores by fitting a statistical model to conflict data, which predicts the likelihood that a dyad will fight as a function of (a) whether the two states are contiguous, (b) the distance between them, and (c) whether either of the states is a "major power." He then uses this model to predict how likely a given dyad is to fight and feeds this likelihood into its "political relevance" score.[16] Finally, to calculate the total number of "politically relevant dyads" in a given year, he adds up each dyad's political relevance score.[17]

[16]For instance, the United States and Mexico are highly politically relevant because the US is a major power and the two countries share a border. On the other hand, Vanuatu and Belgium are hardly politically relevant at all because neither is a major power, and they are also very distant.

[17]Braumoeller's approach, which is developed in Braumoeller and Carson (2011), is essentially a generalization of a binary standard for "political relevance" that has been used in earlier work.

There are also objections that can be levied against this approach to operationalizing the rate of conflict. An ideal metric might would account for the fact that conflicts are not actually independent[18] and for the fact that, intuitively, the "political relevance" of a dyad can change substantially even without changes in any of three input variables.[19]

However, as a baseline, "politically relevant dyads" is a useful and relatively simple metric we can use as our denominator. It probably succeeds in at least *roughly* screening off the effect that changes in the number of states have had on the frequency of interstate conflict. For that reason, **we will follow Braumoeller (2019) in dividing the number of conflicts in a given year by the number of politically relevant dyads.**[20]

3.6.2.3 Summary

We have seen that operationalizing the "rate of conflict" into a specific metric is surprisingly difficult. Every conceivable metric has one important issue or another.

However, for the purposes of trend extrapolation, Braumoeller's choice to focus on **"the rate of interstate conflicts involving the reciprocated use of force, per politically relevant dyad"** is a strong one. We will focus on this metric for the remainder of the paper.

[18]One concern is that it still relies on the (presumably false) assumption that wars are independent. If State A and State B begin to fight, then this will tend to have implications for the likelihood that State A and State C will fight as well. On the one hand, since State A now has less spare "capacity" to fight, it should be more motivated to avoid conflict with State C. On the other hand, if there is a friendly relationship between State B and State C, the initial conflict may provoke a follow-on conflict. The net effect is unclear. If the pacifying effect is stronger than the aggravating effect, on average, then we should expect the number of conflicts to grow *more* quickly than the number of politically relevant dyads. If the pacifying effect is weaker, however, then we should expect the number of conflicts to grow *less* quickly.

[19]As an extreme illustration: the metric would imply that Ancient Rome and China were more politically relevant than the United States and modern China are today. However, it seems strange to count the failure of Ancient Rome and China to go to war as a reflection of peacefulness; they were simply outside of each others' political spheres and could interact with great difficulty. Less extreme versions of this problem arise as we look back into the nineteenth century.

[20]The next best alternative, which may be appropriate for longer-run extrapolation, may be to simply focus on conflicts between contiguous states and use "contiguous dyads" as the denominator.

3.6.3 Choosing a Dataset

Once we have chosen a metric to quantify the "rate of conflict," we need to choose a conflict dataset to apply the metric to. Fortunately, this choice is easy: there are not many datasets to choose from. Unfortunately, we will see, all available datasets are incomplete in ways that make it difficult to identify true longrun trends.

3.6.3.1 The Militarized Interstate Disputes Dataset

The Correlates of War Militarized Interstate Disputes dataset is the most obvious choice for a dataset to use (Palmer et al., 2022). It is the most commonly used interstate conflict dataset in International Relations and has unusually broad and comprehensive coverage. It aims to include all cases of militarized interstate disputes (between "states in the international system") since 1816.

However, the dataset has two important limitations. First, relative to typical notions of what counts as a "state," its early years include only a very incomplete and unrepresentative sample of states. There is a meaningful chance that gradual evolution of its sample – which initially focuses on European states and gradually expands outward to include other states – will create illusory trends in the rate of conflict. Second, although 1816 is of course a fairly early start date, there is also reason to think that trends may look different if the dataset extended further backward into the past.

Let us discuss the sample issue first. For 1816, due to the dataset's Euro-centric inclusion criteria, it includes only twenty-three states.[21] Aside from the United States and Turkey, all of the states are European. Of course, in the early nineteenth century, it is sometimes hard to say exactly which political entities should be considered "states." However,

[21] Specifically, the Correlates of War project criteria for what counts as a "state in the international system" are: "Prior to 1920, the entity must have had a population greater than 500,000 and have had diplomatic missions at or above the rank of chargÃ© d'affaires with Britain and France. After 1920, the entity must be a member of the League of Nations or the United Nations, or have a population greater than 500,000 and receive diplomatic missions from two major powers." The pre-1920 diplomatic mission requirement, in particular, immediately screens off all states that are not politically active within Western Europe.

even if we apply fairly strict standards for statehood, there are clearly omissions. China, for instance, does not appear in the dataset until 1860. This means, for example, that neither of the Opium Wars it fought with European Powers is included.

According to Griffiths and Butcher (2020), who use an inclusive although also intuitively reasonable standard, there were in fact 211 states in 1816.[22] From this perspective, the early years of the Correlates of War dataset include *only about 10%* of states in existence at the time.

This subset of states is also an unrepresentative sample. Europe in the nineteenth century was, in a wide range of ways, very different from the rest of the world. European countries were more powerful, more technologically developed, and more engaged in colonial activity. European culture, political systems, and geography were also quite distinct. It would not be surprising if the rate of conflict was different in Europe than it was in other places.

In fact, at least from qualitative impressions, Europe does seem to have been more peaceful (internally) than many other parts of the world. Nineteenth century European commentators often remarked on the unusual harmoniousness of this era. For instance, writing shortly after the Crimean War, the prominent British historian Henry Buckell (1857) remarked:

> *If we compare one country with another, we shall find that for a very long period wars have been becoming less frequent; and now so clearly is the movement marked, that, until the late commencement of hostilities, we had remained at peace for nearly forty years: a circumstance unparalleled, not only in our own country, but also in the annals of every other country which has been important enough to play a leading part in the affairs of the world.*

[22] Specifically, they define a state as any entity that has: (1) "A population of at least 10,000," (2) "autonomy over a specific territory," and (3) "sovereignty that is either uncontested or acknowledged by the relevant." international actors." The most contestable aspect of this definition is the low population threshold. As a concrete point of reference, this approximately corresponds to the population levels of the very smallest island nations that exist today (e.g. Nauru) or perhaps to a minor Ancient Greek city-state.

Outside of Europe, at the same time, there was a tremendous amount of regional and colonial conflict in other places. It is conceivable, for example, that Europe in this period was multiple times more peaceful than Africa.

The upshot is that the early years of the Correlates of War may meaningfully understate the global rate of interstate conflict. This problem could then create an illusory upward trend — or prevent the detection of a genuine downward trend — as the dataset gradually transitions from (in practice) a dataset of inter-European conflicts to a more comprehensive dataset of interstate conflicts.[23]

Ideally, in addition to including a more complete sample of states, the dataset would also extend back further than 1816. As noted above, 1816 stands out as the first year of a decades-long period that many European commentators considered unprecedentedly peaceful. It also followed three decades of intense bloodshed, spread across the French Revolutionary Wars and Napoleonic Wars. More broadly, the nineteenth century was a notable transitional period in Europe which brought lasting transformations to the character of conflict and to the economies, political systems, and cultures of European states. It is therefore plausible that long-run trends would look fairly different if the dataset were to extend back (for instance) to the beginning of the eighteenth or seventeenth century.

3.6.3.2 More Comprehensive and Longer-Run Datasets

Ultimately, there are concerns about the Correlates of War Militarized Interstate Disputes dataset: (a) its geographic scope is inconsistent over time, (b) its earlier years

[23]There is, on the other hand, something to be said for focusing on trends only within samples of states that are sufficiently interconnected. For example, if we somehow had reliable conflict data for the premodern era, then it would seem strange to look for trends in both European conflict and South American conflict at the same time. The risks of European states fighting would be highly correlated, and the risks of South American states fighting would be highly correlated, but there would presumably be no correlation across the two sets. The fact that the two sets of states are disconnected could be expected to limit how high spikes in the overall rate of conflict can become, unless both sets of states purely by coincidence descend into system-wide conflicts at the same time. Trying to focus our attention only on states that are part of the same "system" (i.e. states whose conflicts are correlated) is obvious one way to limit this spike-dampening phenomenon, as we look further back in time and the world becomes less interconnected. However, even if we do want to take this approach, it still seems that the Militarized Interstate Disputes dataset excludes a more-than-ideal number of states.

include only a partial and non-representative sample of states, and (c) its coverage only extends back to 1816, which was also the start of an anomalously peaceful period.

The first issue is, in principle, easy enough to deal with. We can simply remove various states from the dataset so that it covers roughly the same area from 1816 onward. However, this would come at a cost of a drastic decrease in the sample of cases to study. It would also make the non-representativeness problem worse.

The non-representativeness issue and time window issue can only be addressed by adding more data points. There are a handful of datasets that attempt to go beyond the Militarized Interstate Disputes dataset and more extensively cover interstate conflicts before the twentieth century. The Correlates of War project also maintains a list of additional "extra-state wars" and "non-state wars," some of which (such as the Opium Wars) could easily be reclassified as interstate wars. Unfortunately, as will be discussed in Appendix A, all of these datasets have issues that substantially limit their suitability for trend analysis. For instance, they often omit information (such as the set of states that existed within a given year) that is necessary to calculate rates or disaggregate conflicts in inconsistent ways. They also in general appear to miss an increasingly large portion of cases (even for states they otherwise include) as they extend further back in time and further outside of Europe.

There continues to be value in work to produce more comprehensive, high-quality, and usable datasets of international conflict. Unfortunately, it is also unclear exactly how much room for improvement there ultimately is. The survival of primary source records reporting conflicts sets a fundamental limit on what can be achieved. These records are less likely to survive the further back in time we go and, for the most part, the further we move outside of Europe. Notably, the increasing availability of records as one moves forward in time will by default bias trend estimates upward.[24]

[24]There are also bias issues related to the way in which primary sources filter into the secondary literature, which scholars tend to rely on when producing longer-run or more global conflict datasets. Historical works tend to provide less and less detail the further back in time they go. The author of a reference work – for instance, a handbook providing the overview of the history of a region – is simply much less likely to mention an individual low-level conflict in 1650 than they are to mention a low-level conflict in 1950, even if they are aware of both disputes.

Ultimately, for the time being, the COW dataset is likely the best option for trend analysis. However, we should remain conscious that – relative to an "ideal dataset" with fully comprehensive coverage – it may produce a fairly distorted picture of longrun trends. It is entirely plausible, for example, that an ideal dataset beginning in 1600 and genuinely including all conflicts throughout the world would suggest a fairly different long-run trend.

3.6.4 Refining Our Forecasting Question

Earlier, in Section 3.3, we defined the Peaceful Future Hypothesis in terms that were still somewhat vague. Now that we have chosen a clearly defined metric and historical dataset to focus on, we are now in a position to make the definition more specific.

Our initial primary operationalization of the question was:

> **Peaceful Future Hypothesis (Average Rate Version):** The *average* rate of interstate conflict between 2050 and 2300 will be at least 90% lower than the *average* rate of conflict between 1850 and 2000.

Our secondary operationalization was:

> **Peaceful Future Hypothesis (Peak Rate Version):** The *peak* rate of interstate conflict between 2050 and 2300 will be at least 90% lower than the *peak* rate of conflict between 1850 and 2000.

These now become:

> **Peaceful Future Hypothesis (Average Rate Version, Specific):** Between 2050 and 2030, the *average* rate of conflicts involving the reciprocated use of force (per politically relevant dyad) will be below .0024.
>
> **Peaceful Future Hypothesis (Peak Rate Version, Specific):** Between 2050 and 2030, the *peak* rate of conflicts involving the reciprocated use of force (per politically relevant dyad) will be below .0445.

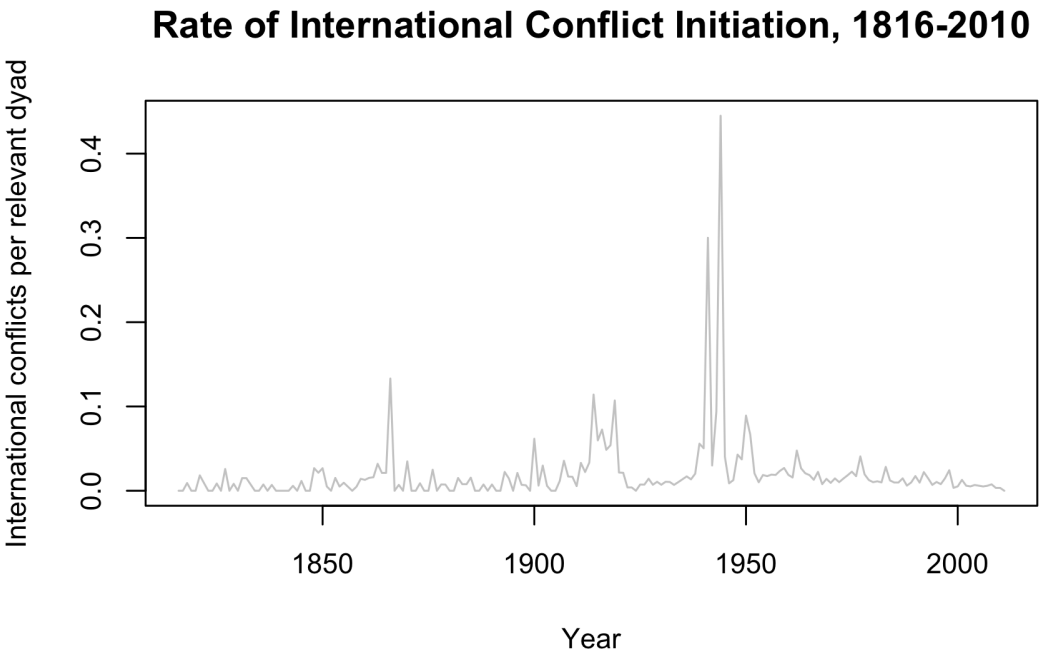

Figure 3.2: The rate of interstate conflict over the past two hundred years, according to the Correlates of War dataset. This particular plot, as well as the following plot, was produced using a slight modification of the replication code for Braumoeller (2019).

These operationalizations, although arbitrary to a large degree, are now specific enough to make forecasting viable.

3.6.5 Identifying and Extrapolating Trends

For the moment, let us put aside any concerns about the reliability or representativeness of our historical data. Suppose that we have settled on a dataset and decided, at least initially, to simply trust it. We have used this dataset to produce a clean plot of the number of our preferred rate-of-war metric, for each year, over a large number of years.

Figure 3.2 shows such a plot. Following Braumoeller (2019), it uses the Correlates of War Militarized Interstate Disputes dataset and the metric "interstate disputes involving the reciprocated use of force, per politically relevant dyad."

The problem of producing a forecast, on the basis of the information in this plot, is a

time-series forecasting problem. In this section, I will briefly outline some general principles for time-series forecasting. I will then apply the principles to the Correlates of War data shown in the figure, to produce an initial "outside view" forecast.

3.6.5.1 Principles of Time Series Forecasting

Time series analysis is the study of trends and patterns in how variables have evolved over time, according to particular datasets. *Time series forecasting* applies the tools of time series analysis to forecast how variables will evolve into the future.

Producing time series forecasts is often a complex process. However, at a high level, time series forecasting tends to involve a few main steps: (i) identifying a statistical model that describes the data well, (ii) fitting the model to the data, and (iii) extrapolating the fitted model out into the future. I will quickly walk through and explain these steps, in the context of a simple example.

In the context of time series analysis and forecasting, a *statistical model* is — essentially — a set of idealized assumptions about how variables of interest evolve over time. For example, consider the following linear model:

$$Y(t) = a + bt + \sigma N(0, 1) \tag{3.1}$$

Here, t is the year and $Y(t)$is the value of the variable of interest in that year. a, b, and σ are unspecified parameters. $N(0, 1)$ is a random variable that follows a normal distribution with mean 0 and variance 1. This model assumes that the variable evolves according to a linear trend — on average, changing by the same amount each year — with normally-distributed random variation around this trend. However – because the model leaves the values of the parameters unspecified – it does not make any assumptions about how high the value is to start (the value of a, how much it tends to increase or decrease each year (the value of b), or the magnitude of the random variation (the value of σ).

Figure 3.3 gives an example of a time series that is obviously well-described by the

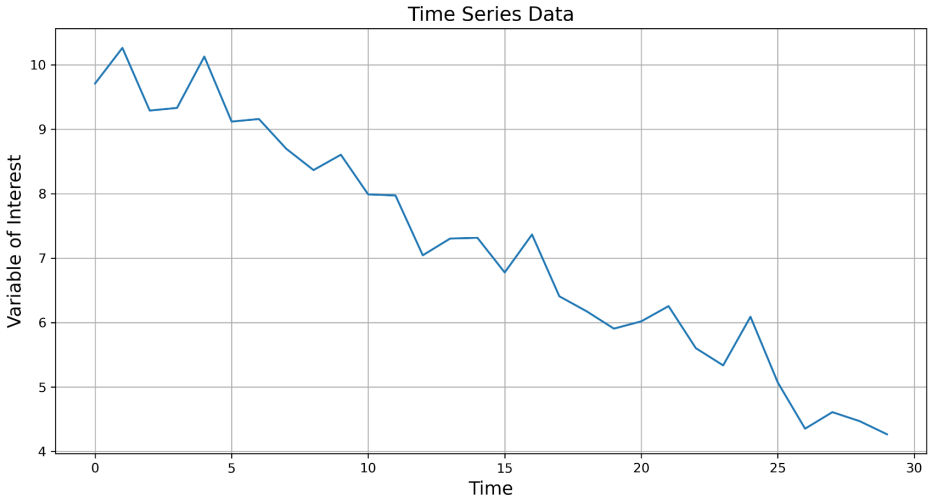

Figure 3.3: An illustrative example of a time series that is well-described by the linear model.

model. Once we notice this kind of match between the model and our data, we can then proceed to the next steps of the process by finding parameter values that produce a good fit. For the data in Figure 3.3, the values $a = 10$, $b = -2$, and $\sigma = -.3$ offer a particularly good fit.

If our fitted model has done a good job describing the variable's evolution for a long time, it is normally reasonable to assume that the model will continue to do a good job describing the variable's evolution for at least *some additional time* into the future.[25] We

[25]This assumption will be discussed extensively below, in Section 3.7.1.2. However, more briefly: One way to justify the assumption is by imagining that, at each time step, there is some unknown risk that the model will "break down" and no longer describe the variable's evolution well. The longer the model goes without "breaking down," the lower this breakdown risk must be. The model goes without breaking down, therefore, the longer we can expect the model to remain valid into the future.

The Laplacian prior – a probability distribution developed by the eighteen century mathematician Pierre-Simon Laplace (Laplace, 2012), in studying cases similar to this one – describes what we should expect if we have no information beyond the fact that a model has described data well for N years so far. On a Laplacian prior, the *expected* time until the next breakdown is N years. The 68% confidence window for the model breakdown is then $.2(N + 1)$ to $5.3(N + 1)$. In the context of a two-hundred-year trend, for example, this implies a confidence window of $[40, 1065]$ years for how long the trend will last.

Of course, in many cases, we have additional information that can make us more or less confident that a given model will break down. For example, we may have additional information about past trend breaks and how frequently they occur. Our model may also imply trends that have implausibly high or low values if extrapolated forward far enough. It also seems likely that the world will change radically enough over the next few hundred years that it would be actively strange if many important trends remained mostly unchanged across this period. Still, the Laplacian prior is useful for setting a rough "anchor point" for how far out we can feel comfortable extrapolating a trend.

can therefore use the model to make predictions. The basic way to do this is to ask: "Suppose the model actually *perfectly* describes how the variable evolves. Then, how likely would the variable be to take on a given value at a given future time?"

We can answer that question by simulating the fitted model several times. We will observe random variation in how the variable evolves, due to the fact the model incorporates randomness. If we run enough of these simulations, though, we can then estimate *how likely* the variable is to take on a certain value by noting *how frequently* it takes on the value.

We can further improve our forecasts by incorporating uncertainty about best-fit parameter values. The natural way to do this is to sample from a statistical distribution of parameter values (weighted by the relative quality of their fit) when running our simulation.

Beyond this point, we can try to account for several deeper forms of uncertainty: the possibility of changes in parameter values (in the past or in the future), the possibility that our dataset is importantly flawed, and the possibility that our model does not (or will no longer) do a good job of describing how the rate of conflict evolves. However, accounting for these deeper sources of uncertainty will be left for the next main section of the paper.

3.6.5.2 Selecting an Initial Model

We can now return to Figure 3.2, which plots the Correlates of War time series.

Something that is immediately obvious is that the simple linear model does not describe the time series well. We are clearly not looking at normal random variation around a simple trendline. The time series display seems to display irregular *cycles*: the average rate of conflict will shift up to a higher level for a time, then shift back down to a lower level for a time. The pattern of year-by-year variation, which often includes sharp spikes, also does not suggest a normal distribution.

The cycles are especially problematic. Braumoeller (2019) identifies several level changes and rules out the possibility that anything like the simple linear model could have

produced them by chance.[26] (See Figure 3.4.) Any model that can accommodate these observations will necessarily be more complex.

One model that may be able to roughly accommodate our dataset is a *geometric Brownian motion model*. This model is often used, for instance, to study the evolution of stock prices.

The way the model works is that it assumes that, in each time period, the variable of interest increases or decreases by some random multiplicative factor. This leads the variable to wander around somewhat erratically. If the random factors happen to be much greater than one a few times in a row, for example, then the variable may rise to a high value and not come back down for a while.

More specifically, the standard geometric Brownian motion model is given by the following equation:

$$Y(t) = Y(t-1) \cdot e^{a+\sigma N(0,1)} \tag{3.2}$$

In this equation, the random factor is given by $e^{a+\sigma N(0,1)}$. $N(0,1)$ is again a normally distributed random variable with mean 0 and variance 1. a is a constant that describes how "biased" the random factor is to be greater than or less than one. If $a = 0$, then, at each time step, the variable of interest is equally likely to increase or to decrease.

To provide some intuition for this model, in the context of conflict, we can imagine that the rate of conflict is tied to the level of "instability" in the international system. The level of instability shifts up or down each year. The exponential nature of the shifts implies that the more unstable the system becomes, the more quickly it can become even more unstable. The more ongoing conflicts there are, for example, the greater the potential is for a system-wide conflict to emerge. There also may or may not be some bias toward either

[26] Actually, it is more correct to say that Brauomoeller rules out the possibility that a *stationary model* could have produced this data. A stationary model is one where the probability that $Y(t)$ will take on a given value is the same in each time period t. However, it is also clear from his analysis that a linear model would not fit the data well either.

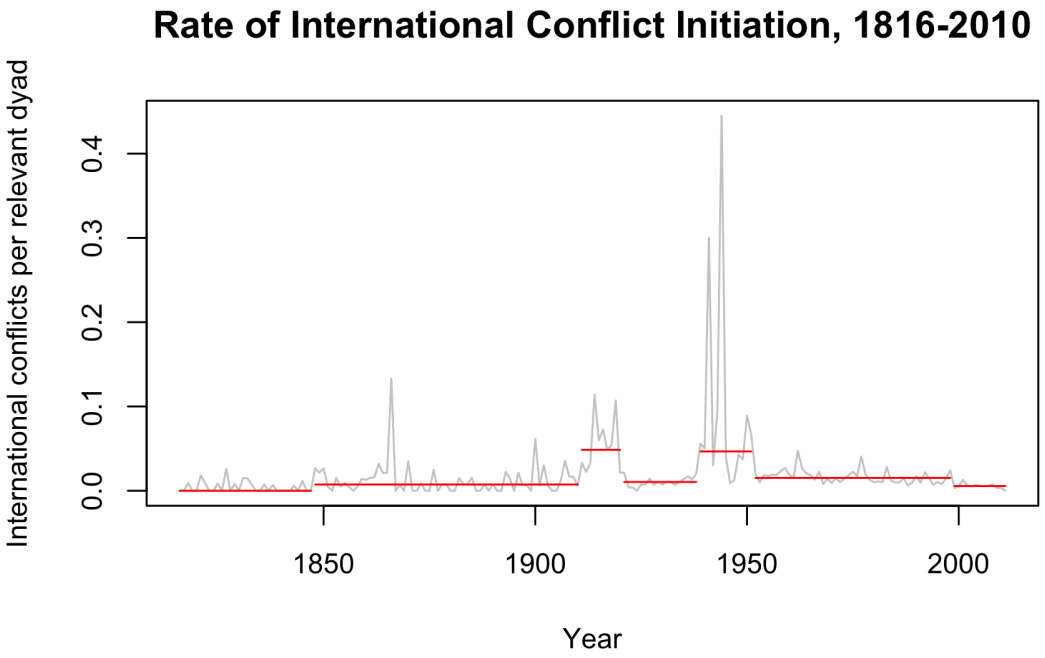

Figure 3.4: This figure is also adapted from Braumoeller (2019). Braumoeller considers whether the time series could have been produced by a *stationary model*, meaning: a model where the probability distribution for the value of $Y(t)$ is the same in each year, rather than changing over time or depending on its value in previous years. He finds that stationary models cannot fit the data well. The breaks in the red lines in this diagram correspond to points where the rate of conflict varies more than it plausibly could by chance, if a stationary model were correct. The heights of the red lines then correspond to the median rate of conflict within this different periods. Since the heights of the lines go up and then down, there is also clearly no way that a simple linear model — even one where the error term does not follow a normal distribution — would be sufficient to produce this pattern. Ultimately, any model that is capable of fitting the data well must be at least somewhat more complex.

upward or downward shifts in instability. For example, a downward bias could reflect the natural tendency for states to try to maintain stability.

Unfortunately, seemingly, there is at least one important problem with this model. A key property of geometric Brownian motion — which does not seem to fit the case of international conflict very well — is that it can in principle wander up to indefinite heights. In practice, the rate of conflict cannot really exceed 1.

One simple way we can remove this undesirable property of the model would be to

simply place a ceiling on the rate of conflict; the rate of conflict stops moving upward once it hits the ceiling. A more elegant way to deal with the problem, though, is to hypothesize an "equilibrating force" that smoothly increases the bias toward de-escalation as the rate of conflict becomes more extreme. The pattern of spikes in Figure 3.2 – specifically, the fact that periods of high conflict are so fleeting – weakly suggests the existence of this kind of equilibrating force. The idea of this equilibrating force is also fairly intuitive: As the world's level of conflict becomes increasingly extreme, states may become increasingly motivated to pursue stability.

We can introduce an equilibrating force by making the following modification to the model:

$$Y(t) = Y(t-1) \cdot e^{a+b\log(Y(t-1))+\sigma N(0,1)} \tag{3.3}$$

Now — so long as the value of b is negative — higher values of Y will result in a stronger bias toward Y decreasing in the next time step.

3.6.5.3 Fitting the Initial Model

How well does this new model match our dataset? To see, we can attempt to find best fit parameters. It will be easiest to do this if we transform the above equation into a linear expression. Calculating the ratio between each datapoint and the preceding datapoint, and then taking the logs of the ratios, gives us the following linear expression:

$$\ln\left(\frac{Y(t)}{Y(t-1)}\right) = a + b\log(Y(t-1)) + \sigma N(0,1) \tag{3.4}$$

Now that we have a linear expression, we can find best-fit parameters by running a standard linear regression. As reported in Table 3.2, the regression yields the values $a = -1.8$, $b = -.4$, and $\sigma = .7$. The value for b is also highly statistically significant

Table 3.1: Regression on conflict rate data, for our initial model

	Dependent variable:
	Log(CurrentRate/PreviousRate)
Log(PreviousRate)	-0.418^{***}
	(0.078)
Constant	-1.767^{***}
	(0.327)
Observations	129
R^2	0.184
Residual Std. Error	0.738 (df = 127)
Note:	*p<0.1; **p<0.05; ***p<0.01

($p = 4 * 10^{-7}$), suggesting that there may in fact be an equilibrating force.[27]

We can now form an initial impression of how well our model matches the data, by simulating the model several times. Looking at several simulations should give us a sense of whether the model's behavior is at least qualitatively similar to what we have observed historically. Figure 3.5 shows four simulated "counterfactual histories" of the period between 1816-2010, the same period covered in the dataset.[28]

The pattern is, at a glance, quite similar.historical data relatively well.

3.6.5.4 Representing Trends

Figure 3.6 — which shows the expected rate of war over time, calculated by taking the average of 10,000 simulations — makes it clear that our fitted model does not display a long-run trend.

The absence of a trend is to be expected. Our model now, largely by design, has a

[27]Note: Approximately a quarter of data points are thrown out, due to the rate of war in that year or the preceding year being zero. The value of $log(Y(t)/Y(t-1))$ is not well-defined in these cases. Throwing out zero years may introduce a degree of bias into our estimates. Hopefully, the bias is not too large, since we are ignoring an equal number of years where the rate of conflict *drops down* to the zero level and years where it *returns* from the zero level. However, we also cannot rule out the possibility of meaningful bias.

[28]I start the simulation — and also subsequent simulations — with an assumed rate of conflict of .01 in the first year. In fact, in 1816, there were no recorded conflicts.

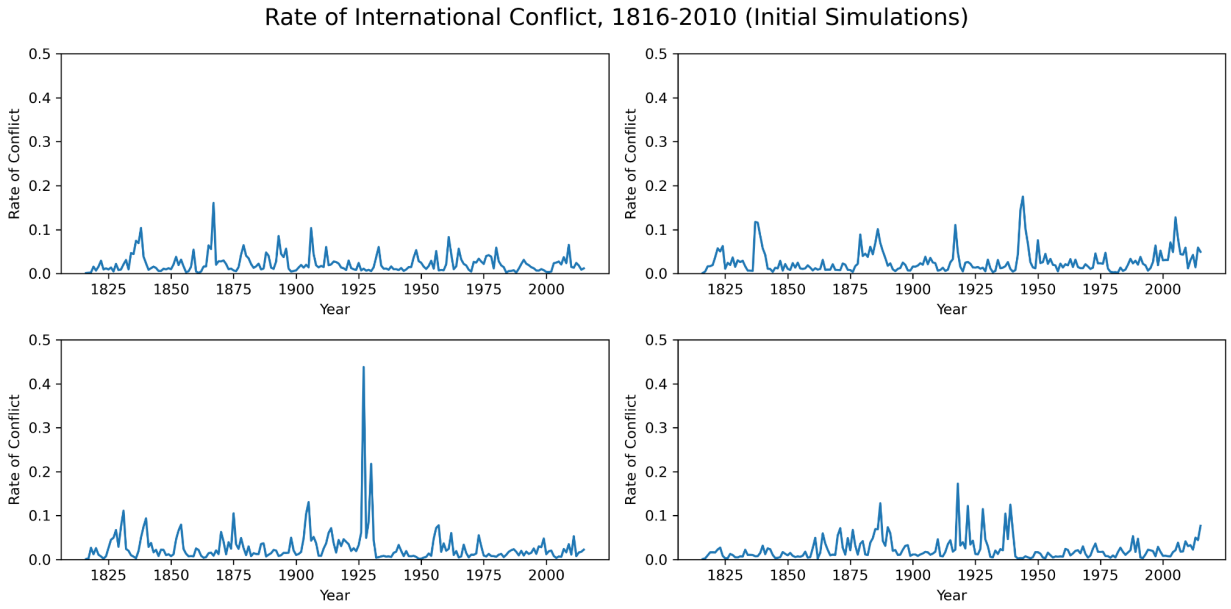

Figure 3.5: Four "counterfactual histories" of the rate of conflict over the past two centuries. These counterfactual histories were produced by fitting our model to historical data from this period, then extrapolating the fitted model forward from the start of the period.

tendency to equilibrate. As the rate of conflict grows, the downward bias grows. As the rate of conflict shrinks, the downward bias shrinks and — at some point, once it becomes negative — actually switches into an upward bias. Although individual rollouts of the model may display rising and falling patterns, they are the result of random noise and do not reflect an underlying long-run trend.

If we want to *allow* for the possibility of an underlying trend, then we need to make one final modification to our model. We can introduce an additional time-dependent bias term, ct. The model now becomes:

$$Y(t) = Y(t-1) \cdot e^{a+b\log(Y(t-1))+ct+\sigma N(0,1)} \tag{3.5}$$

If the value of c is negative, for example, then the downward bias in Y's movements becomes increasingly strong as time passes. More concretely, we might imagine that — due to some growing pacifying force, such as increasingly strong economic incentives or moral aversion to war— states become increasingly able or inclined to maintain stable relations

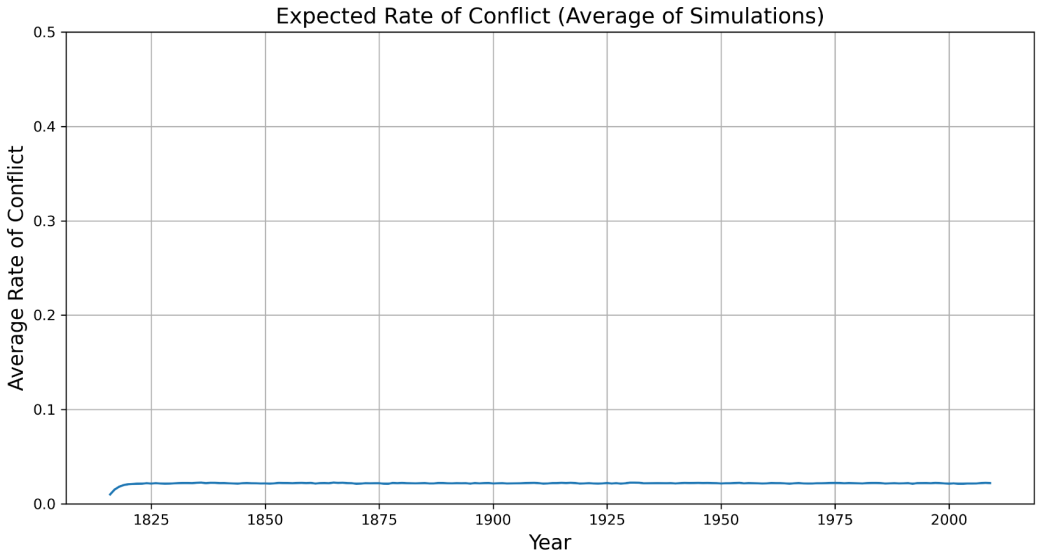

Figure 3.6: The expected rate of conflict over the past two centuries, according to our fitted model. The expected rate is estimated by simulating our model many times and then calculating the average rate across these simulations. Notably, but unsurprisingly, there is no observed trend.

over time. Therefore, over time, increases in instability become smaller and less common. When instability does rise, states also react in ways that are less likely to cause instability to grow further.

In this new model, there is still a tendency toward equilibrium. However, if the value of c is non-zero, *the equilibrium point will shift upward or downward over time*. It is now possible for the model to display an underlying long-term trend.

3.6.5.5 Fitting the Revised Model

Finally, we can fit this revised model to our data and see if the fitted parameters imply a trend. Table 3.2 reports the results of another linear regression.

We find no evidence of a trend. The value of the trend parameter is negative (-.0006), but its p-value (.7) is very far from the statistically significant level.

Of course, at the same time, this finding is consistent with the existence of a trend that is small enough to fall within the confidence interval around the point estimate. If we

Table 3.2: Regression on conflict rate data, for version of our with a trend term

	Dependent variable:
	Log(CurrentRate/Previous)
YearsSince1816	−0.001
	(0.001)
Log(PreviousRate)	−0.423***
	(0.079)
Constant	−1.717***
	(0.350)
Observations	129
R^2	0.185
Residual Std. Error	0.741 (df = 126)
Note:	*p<0.1; **p<0.05; ***p<0.01

consider a 95% confidence interval, then anything in the interval $[-.002, .003]$ is consistent with the data.

In practical terms, this implies a range of possible trends. Figure 3.7 shows the trendlines implied by the lower, central, and upper estimates of the trend parameter.[29] The upper bound implies that the expected rate of conflict has increased by about a *factor of three* over the past two centuries. The lower bound implies that the expected rate of conflict has fallen by about a *factor of four*.

Overall, these regression results suggest that our data is not consistent with the existence of a *radical* upward or downward trend in the rate of conflict. However, it is still consistent with the existence of a *substantial* trend. The range of possible trends is still fairly broad.

[29] Assuming specific values for the trend parameter also has implications for the point estimates of the other parameters. In turn, the values of these other parameters have implications for the initial "equilibrium point" of the rate of conflict. This is why the three sloping lines in the graph begin at different heights.

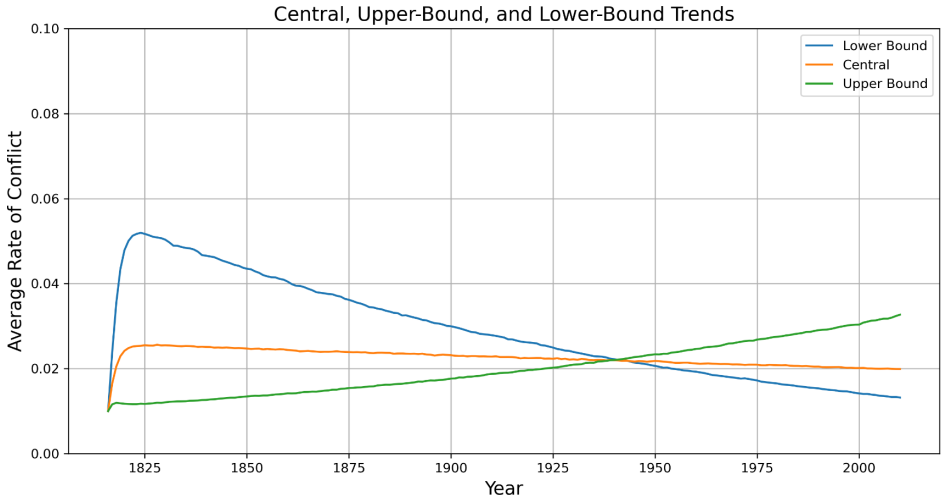

Figure 3.7: The expected rate of conflict over the past two centuries, estimated using central, lower-bound, and upper-bound estimates of the time trend parameter. We can see that a wide range of different underlying trends could have produced the pattern of conflict we observed over this period.

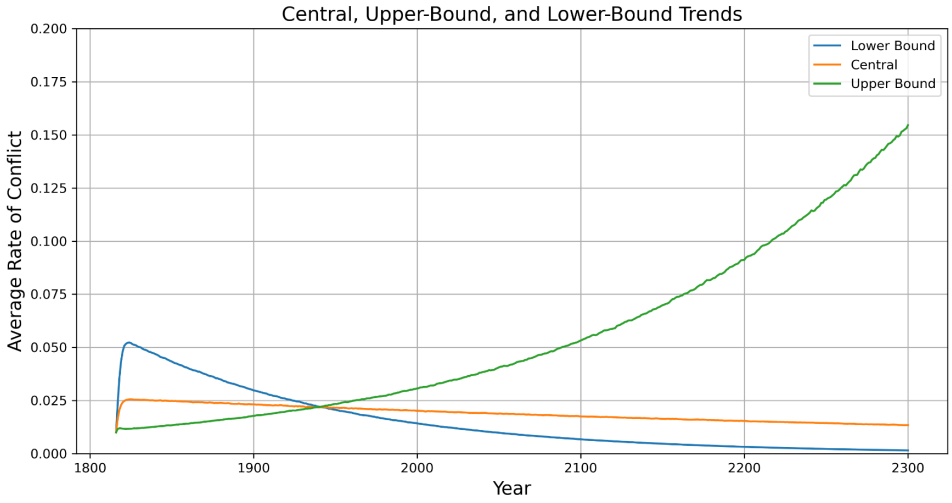

Figure 3.8: The expected rate of conflict up until 2300, estimated by continuing to extrapolate the model forward.

3.6.5.6 Extrapolation

What do we find if we extrapolate the models forward?

Figure 3.8 extrapolates the trendlines forward all the way up to 2300. The upper bound estimate implies that, start-to-end, the expected rate of conflict increases by about a *factor of 15*. The lower bound estimate implies that, start-to-end, the expected rate of conflict decreases by about a *factor of 33*.

Of course, trend graphs present an artificially neat picture. The randomness inherent in our model implies that a very wide range of trajectories are consistent with any given underlying trend. Even if there is a slight underlying downward trend, for example, there may still be sharp upward spikes in the rate of conflict, or the rate of conflict may drift upwards for an extended period purely by chance.

Figure 3.9 gives a clearer sense of the diverse range of possibilities, by plotting several possible trajectories for the period 2024 to 2300.[30] The simulations were generated by randomly sampling different sets of parameter values, weighted in proportion to how likely they would have been to produce our historical time series data.

We can also, now, generate an overall forecast of the likelihood of a peaceful future. Recall that we have operationalized the Peaceful Future Hypothesis in the following ways:

- **Peaceful Future Hypothesis (Main Operationalization):** The *average* rate of conflict between 2050 and 2300 will be at least 90% lower than the *average* rate of conflict between 1850 and 2000.

- **Peaceful Future Hypothesis (Secondary Operationalization):** The *peak* rate of conflict between 2050 and 2300 will be at least 90% lower than the *peak* rate of conflict between 1850 and 2000.

More specifically, given the rate metric and data we are using, the first operationalization implies that the *average* rate of conflict must be below .0024. The second opera-

[30] I have also arbitrarily assumed a rate of conflict of .001 for 2024.

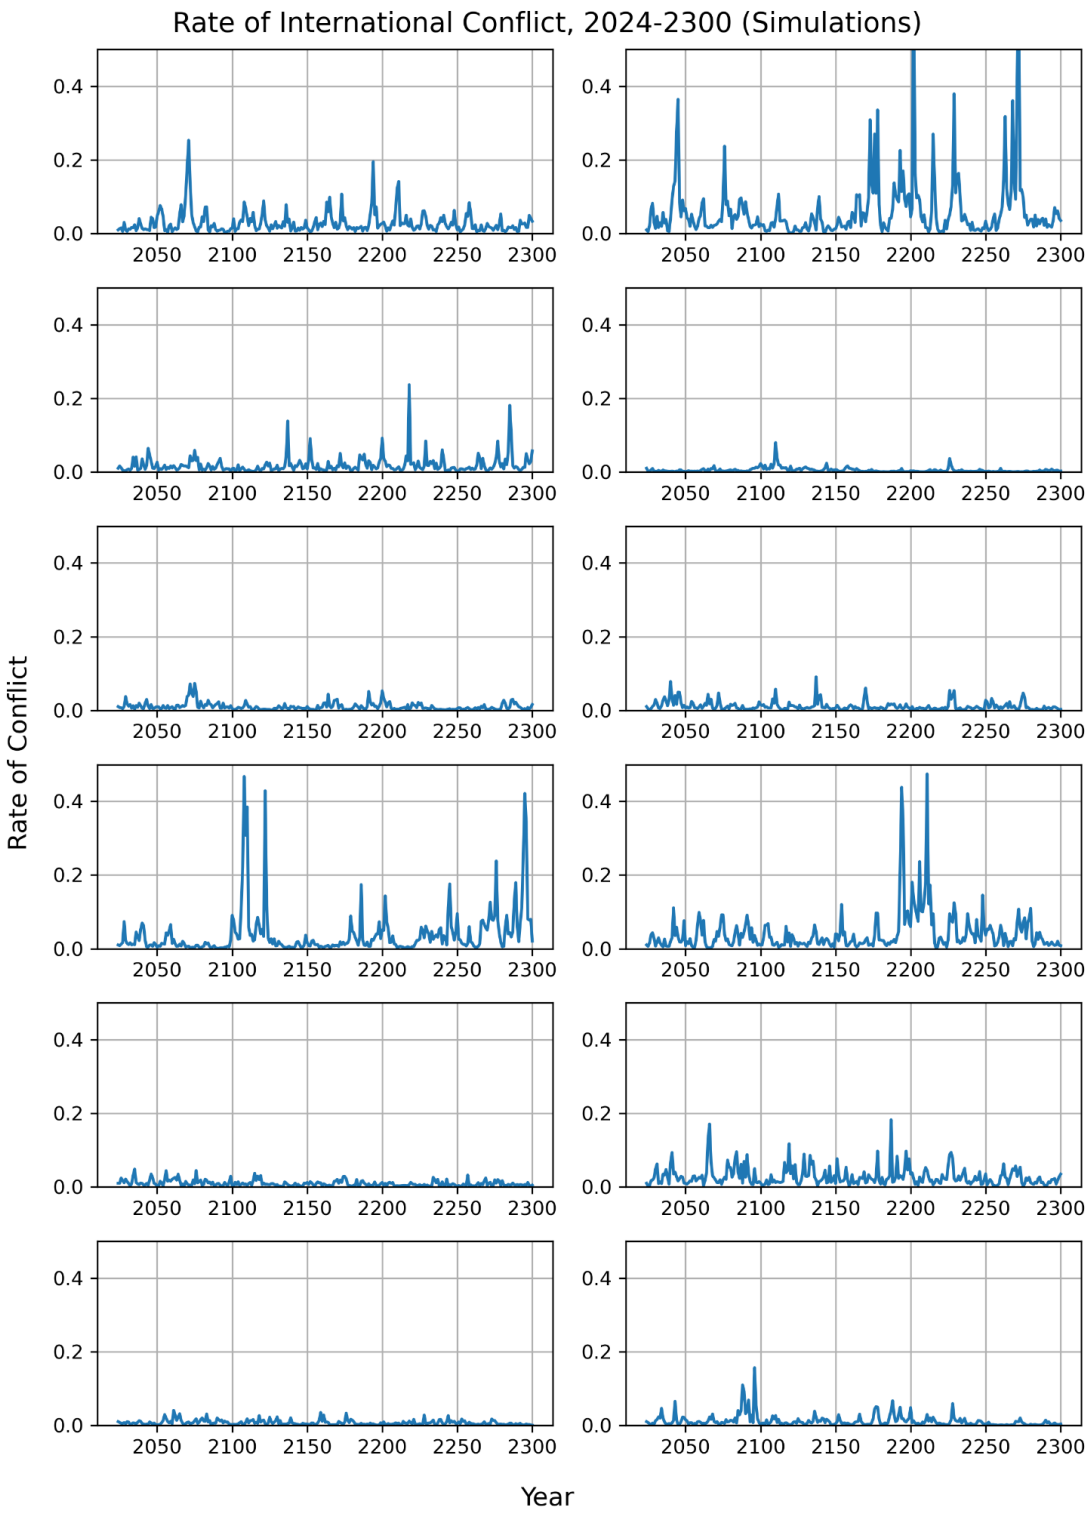

Figure 3.9: Twelve simulations of the next two centuries of conflict. These simulations were produced by fitting our model to historical data, then extrapolating it forward multiple times. To account for parameter uncertainty, we sample from a distribution over the parameter values.

tionalization implies that the *peak* rate of conflict must be below .0445.

To estimate the probabilities of these outcomes occurring, we can simulate the model many times — just like we did to produce Figure 3.5 — and calculate how frequently the average and peak rates fall below the relevant threshold. As above, to account for parameter uncertainty, we can randomly sample different sets of parameters each time.

Across 100,000 simulations, we find that that *average* rate falls below the relevant threshold only **1%** of the time. Meanwhile, the *peak* rate falls below the relevant threshold **7%** of the time.[31]

Expressed probabilistically — and using our main operationalization — **our model implies that there is only a 1% chance of a peaceful future.**[32] According to our model: Even though there may be a meaningful downward trend in conflict, and even though there is a meaningful chance that the future will be substantially more peaceful than the past, it is extremely unlikely that the future will be *radically* more peaceful than the past.

Initial Outside View Probability Estimate: 1%

3.7 Various Complications: Modifying Our Initial "Outside View" Forecast

We have now produced an outside view forecast, based purely on time series data. It implies that — given our specific criteria for what counts as a "peaceful future" — the future is very unlikely to be peaceful. Specifically, it assigns only an approximately 1% probability to a peaceful future.

[31] As a point of comparison, if we rerun these simulations with the trend parameter to zero, then the probabilities become effectively zero. These probabilities are driven by the possibility of a meaningful downward trend.

[32] Less strict operationalizations yield higher probabilities. For instance, if we switch two using 80% thresholds rather than 20% thresholds, then our primary and secondary operationalizations are satisfied 5% of the time and 31% of the time respectively.

However, our time series forecast is based on three key assumptions:

1) **No structural breaks:** The "true values" of the model's parameters (especially the trend parameter c) have not changed in the past two centuries and will not change in the next two either.
2) **Dataset accuracy:** The historical conflict dataset that fitted our model to is accurate and unbiased.
3) **Model accuracy:** The process by which the rate of war has evolved (and will continue to evolve) is well-described by our model.

All of these assumptions are highly questionable. We know that no trend can last forever and that even longstanding trends must break at some point. We know that our model is a vast simplification of a complex and mysterious real-world phenomenon. Finally, as discussed in Section 3.6.3.1, we know that our data only describes the rate of conflict within an incomplete and geographically inconsistent set of states.

Our overall outside view forecast should in some way account for these additional sources of uncertainty. This section of the paper is dedicated to clarifying how we can account for them.

One recurring theme is that accounting for these deeper forms of uncertainty often requires making *judgment calls*, where a wide range of decisions could be considered reasonable. At the same time, clearly, some decisions are more reasonable than others; this prevents the probability estimate that we arrive at from being entirely arbitrary.

The most significant finding in this section is that: accounting for possible trend breaks — both past and present — can easily make a drastic difference to our forecast. Accounting for the possibility of past trend breaks initially leads to a dramatically higher probability estimate. There is, in fact, statistical evidence that the rate of conflict has begun to decline since the mid-20th century. Extrapolating this recent trend forward then suggests a *very high* likelihood of a peaceful future. However, accounting for the possibility of additional

future trend breaks warrants adjusting the optimistic estimate back down. The question of how we should make these consequential adjustments is both difficult and requires making a number of judgment calls.

On the other hand, this section will also argue that data uncertainty is less of a problem than it may initially seem. One reason for this is that: if the strongest case for a peaceful future rests heavily on extrapolating recent trends, then uncertainty about data that substantially predates these trends has only limited relevance. Overall, I suggest that there is no need to substantially adjust our forecast to account for data uncertainty.

Model uncertainty is a more significant issue. In this case, I suggest that model uncertainty should overall lead us to adjust our probability estimate slightly downward. The most important consideration is that we should also give meaningful weight to the simpler version of the model that lacks the trend term; in part, this is a way of acknowledging that the null hypothesis ("no trend") deserves special *a priori* weight.

Along the way, in this section, I will track how the overall probability estimate evolves as each additional complication is considered. I will ultimately arrive at **a final outside view probability estimate of 20%**, although meaningfully higher or lower estimates could also be defended.

3.7.1 Structural Breaks

Even if we assume that our model is accurate, we may still produce misleading forecasts if we fail to account for *structural breaks*. These are possible points where the values of the model's parameters change.

Most notably, our forecast is based on the assumptions that (i) there has been a consistent trend (i.e. value of $\boldsymbol{c}$) for the past two centuries and that (ii) the trend will continue for the next two centuries as well. It is far from obvious, though, that this assumption is true.

3.7.1.1 Past Structural Breaks

3.7.1.1.1 Detecting Past Structural Breaks If there have been major structural breaks within the period covered by our dataset, then our parameter estimates may be inaccurate.

For instance, as an illustration, suppose there has been a structural break that divides the dataset into two distinct periods. In the earlier period, there is a substantial upward trend. In the current period, there is a substantial downward trend. If we try to fit our model to the full dataset, in this scenario, then we could wrongly conclude that there has been no substantial trend. We would fail to recognize that there actually currently is a downward trend.

In the context of our time series model, we are looking for evidence that the values of our parameters — especially c — have changed at some point. The Bai-Perron test is one common statistical test for detecting these kinds of changes. This test uses the "Bayesian Information Criterion" (BIC) as a metric to compare the predictive power of variants of a model with different numbers of breaks, while penalizing models with larger numbers of breaks.[33]

Lower BIC values indicate that a model is more likely to be "correct." More specifically, if a pair of models have BIC values BIC1 and BIC2, then — if we make certain statistical assumptions, and if we have no *a priori* reason to consider one model more likely than the other — then the relative probabilities of the two models are given by:[34]

[33]In this paper, my application of the Bayesian Information Criterion — to assign probabilities to different models and then ultimately use these probabilities to produce weighted forecasts — broadly follows the approach laid out in Wasserman (2000). Although the criterion is a commonly used statistical tool (Neath and Cavanaugh, 2012), there are also some valid concerns about its use. One of these concerns is that its use is grounded in the idealized assumption that one of the models being examined is actually true, or at least, as Neath and Cavanaugh put it, "*quasi-true.*" Gelman et al. (2014) have expressed principled concerns with its use partly for this reason. Weakliem (1999) has presented a number of additional critiques, including the critique that the Bayesian prior probabilities it implicitly uses often diverge from the priors it would in fact seem reasonable to hold.

[34]As a few points of reference: a difference in BIC values of 1 implies that one model is about *1.6 times* more likely to be correct, difference of 2 implies that it is about *2.7 times* more likely, a difference of 6 implies that it is about *20 times* as likely, and a difference of 10 implies that it is about *148 times* as likely. A common rule of thumb is to treat BIC differences below 2 as essentially meaningless for discriminating

Table 3.3: Results of Bai-Perron test, assessing the relative likelihoods of different numbers of past trend breaks.

Number of Past Breaks	**Most Likely Break Years**	**BIC score**	**Relative Probability**
No Past Breaks		305.1	11%
One Past Break	1937	302.4	41%
Two Past Breaks	1921, 1944	302.5	39%
Three Past Breaks	1894, 1921, 1944	305.5	9%
Four Past Breaks	1894, 1921, 1944, 1989	320	0%
Five Past Breaks	1894, 1921, 1944, 1963, 1989	338	0%

$$\frac{\Pr(\text{Model 1})}{\Pr(\text{Model 2})} = e^{-\frac{1}{2}(\text{BIC1 - BIC2})} \tag{3.6}$$

Now, running the test, *we find that the BIC value is lowest (BIC = 302.4) for a variant of the model with a single trend break occurring in 1937.* If we also perform a Chow test — a common test for judging whether a model's parameter values differ before and after a *particular* point — then we again receive an indication that the model's parameter values are different before and after 1937. The test detects a change with a very high level of statistical significance (p = 3.9E-7).

The two-break variant of the model — with one break in 1921 and another in 1944 — also has essentially an equivalently low BIC score. The BIC is also only moderately higher for the no-break and three-break variants of the model. Table 3.3 presents the full results of the Bai-Perron test.

Let us focus, for now, on the scenario that the test suggests is most likely, where there has been a single mid-twentieth-century trend break.

Table 3.4 reports the results of three relevant regressions. The first regression is fitted to the full dataset. The second regression is fitted only to data points *before* 1937. The third

between models and BIC values above 6 as highly meaningful.

Table 3.4: Regressions on data subsets

	Dependent variable:		
	Log(CurrentRate/PreviousRate)		
	(1)	(2)	(3)
	All Years	Year<1937	Year>1937
Year	−0.001	0.002	−0.026***
	(0.001)	(0.003)	(0.005)
Log(PreviousRate)	−0.423***	−0.494***	−0.805***
	(0.079)	(0.133)	(0.117)
Constant	−0.668	−5.846	47.626***
	(2.705)	(6.085)	(9.458)
Observations	129	57	72
R^2	0.185	0.207	0.409
Note:			*p<0.1; **p<0.05; ***p<0.01

regression is fitted only to data points *after* 1937.[35]

Notably, the time coefficient in the third regression is highly statistically significant. *This strongly indicates a downward trend in the rate of conflict in the post-1937 period.*

3.7.1.1.2 Extrapolating More Recent Trends Forward Let us now extrapolate our model forward, using the new parameters we have found for the post-1937 period. Figure 3.10 presents four simulations that start in 1937 and run until 2300. As above, these simulations were produced using random sampling from the distribution of parameter values.

Notably, these plots all show that the rate of war rapidly declining. There are occasional spikes and periods with upward drift — for example, in the third model, conflict appears to be on an upward trajectory 1960 to 1980 — but in cases the rate of conflict converges on a very low rate by the end of the 21st century.

[35]The parameter values of the third regression do not change much if we instead start it in 1944, in line with two-break and three-break models.

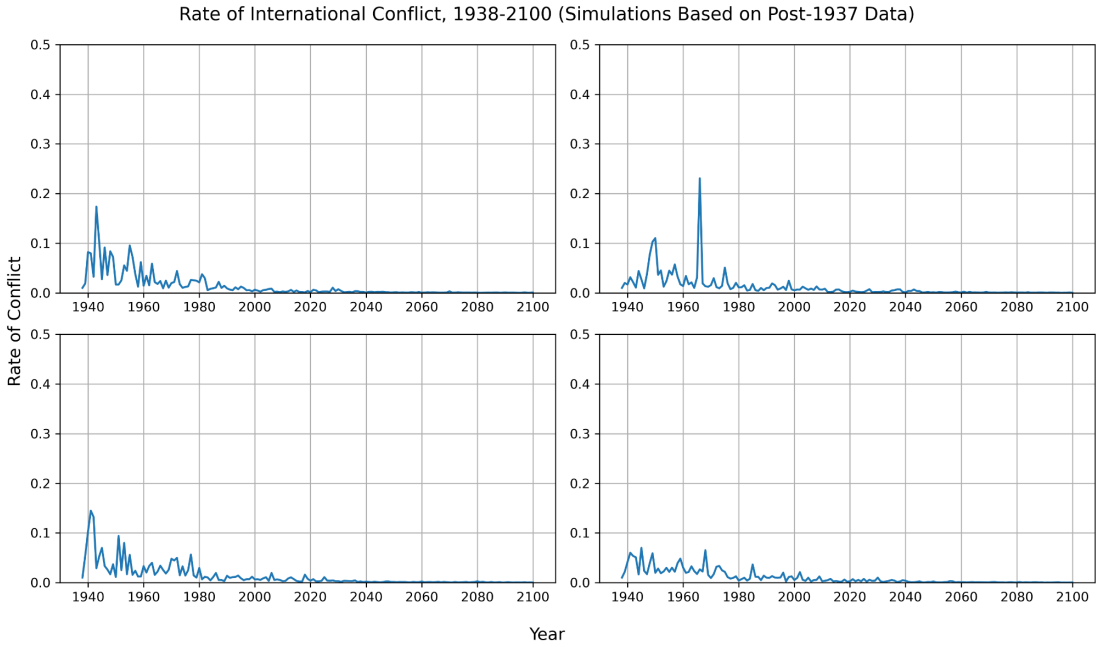

Figure 3.10: Four simulations of the rate of conflict from 1937 onward, generated by fitting the model only to post-1937 data. Random sampling has been used to select the parameter values.

We can now use our new parameter distribution to assess the likelihood of a peaceful future. By running the simulation until 2300, we can calculate how frequently the average rate of conflict between 2050 and 2300 falls below .0024. We find that this occurs **98%** of the time. We also find that the peak rate of conflict in this period falls below .0445 approximately **100%** of the time.

In other words, *under certain assumptions* — namely, the assumptions (i) that our underlying model is valid and the data we are using is accurate and unbiased, (ii) that there will be no further trend breaks before 2300, and (iii) that the most recent trend break was in 1937 — *the future appears virtually certain to be peaceful.* Of course, as already noted, these are all questionable assumptions.

For a moment, let us hold onto the first two assumptions while dropping the third. What should we predict, now that we recognize that there *may* have been trend breaks over the past two centuries?

The approach we will follow is to assign relative probabilities to variants of our model

Table 3.5: Calculation for the likelihood of a peaceful future, given the simplifying assumption that there will be no future trend breaks.

Number of Past Breaks	**Probability This Many Breaks Have Happened**	**Probability of Future Peace, Assuming No Future Breaks**	**Contribution to Total Probability of Future Peace (Product of Columns)**
No Past Breaks	11%	1%	0%
One Past Break	41%	98%	40%
Two Past Breaks	39%	87%	34%
Three Past Breaks	9%	87%	8%
Total Probability of Future Peace (Sum of Rows)			82%

with different numbers of breaks, produce different forecasts using these different variants, and then produce a combined forecast by averaging across them. Table 3.5 presents the calculation, focusing on the "average rate" operationalization of the Peaceful Future Hypothesis.

Our new composite forecast implies that there is an approximately **80% chance** of a peaceful future.[36] This is, of course, a major increase from our previous forecast of 1%.

Current Outside View Probability Estimate: 80%

3.7.1.2 Future Structural Breaks

Of course, if the trend in conflict may have changed in the past, then this suggests it may change in the future as well. Even if it has not changed over the past two centuries, we know that trends must always come to an end eventually. At some point, any trend will begin to run up against a wall, or something fairly fundamental in the world will shift and disrupt the trend.

Forecasting future structural breaks is a difficult problem. One way to get an initial handle on the problem is to estimate the annual likelihood of a structural break, based on

[36]I will tend to round significant digits in estimates, in order to avoid any false implication of precision.

how frequently they seem to have occurred in the past.

Following an estimation approach pioneered by the eighteenth century mathematician Pierre-Simon Laplace (2012), we can model trend breaks as random events that have an equal chance of occurring in each time period. This means that, by assumption, each year there is an equal chance that the current trend in the rate of war will change. It is as though a weighted coin is being flipped each year, and, if the coin happens to land tails, then the current trend breaks. The lower the annual odds of a trend break are — by analogy, the more heavily weighted the coin is — the longer the trend will tend to last.

The Laplacian approach is to start out by assigning each possible annual likelihood (i.e. weighting of the coin) an equal probability. We then update away from this uniform prior probability distribution by observing how frequently trend breaks occur. Once we have a new posterior probability distribution, regarding the annual likelihood of breaks, we can use this to calculate a distribution over the *length of time until the next break*.

There is no simple analytic expression for this probability distribution, but we can estimate it numerically. Figure 3.11 shows the likelihood that the current trend will persist for various lengths of time, under different assumptions about how many trend breaks have already occurred.

We can see that, even if we assume there have been one or two trend breaks so far, there is at least a *non-negligible* chance that the current trend will persist into the next century.

Next, we need to try to use this information to update our overall estimate of the probability that the future will be peaceful.

One crude approach is to initially assume that any future trend break will lead to greater conflict. This initial assumption — although it inevitably biases our estimate downward — will simplify our analysis. Let us go further and make these two more specific assumptions:

- **Early Breaks Rule Out Peaceful Futures:** If the current trend (whatever it is)

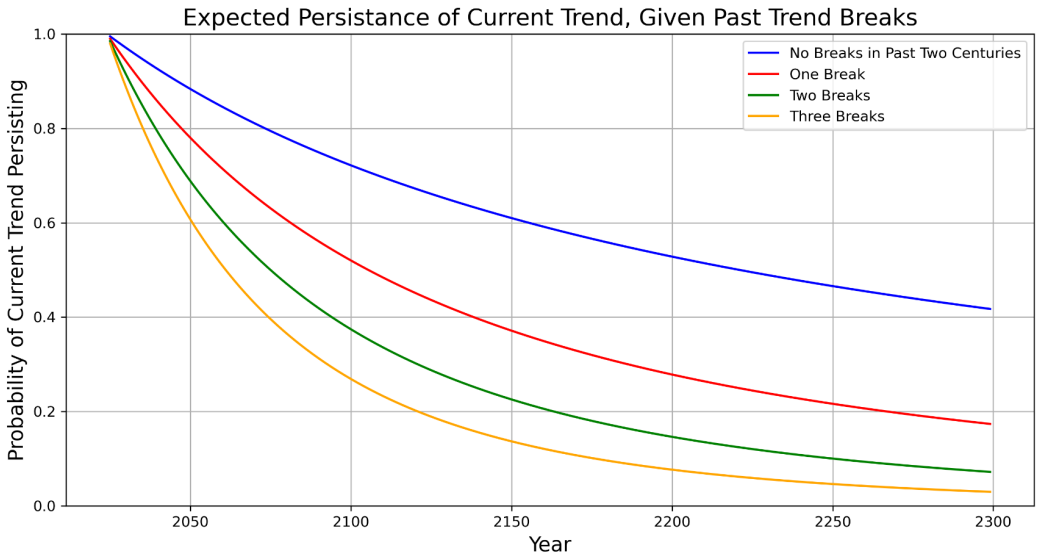

Figure 3.11: The fewer trend breaks we believe there have been since 1816, the longer we should expect the current trend to last.

breaks before 2150, then the period between 2050 and 2300 will not qualify as peaceful.

- **Late Breaks Can Be Ignored:** If the current trend (whatever it is) lasts until at least 2150, then this trend break will only have a negligible influence on the likelihood that the period will qualify as peaceful.

Given these two assumptions, we can now generate an overall probability estimate of the likelihood of a peaceful future. Table 3.6 represents the relevant calculation.

Our new overall probability estimate is that there is a **15%** chance of a peaceful future.

This should be understood as a *conservative lower bound*, though. We have essentially assumed that trend breaks can only be sharply negative. Of course, trend breaks can also push in the direction of greater peace or cause only minor changes in direction. Our downward-biased approach of accounting for future trend breaks has dropped the probability all the way from 82% to 15%. A more balanced estimation approach would presumably lead to a smaller drop.

Table 3.6: Calculation for the likelihood of a peaceful future, accounting for the possibility of future trend breaks.

Number of Past Breaks	**Probability This Many Breaks Have Happened**	**Probability of No Future Break Before 2150**	**Probability of Future Peace, Conditional on No Future Break**	**Contribution to Total Probability of Future Peace (Product of Columns)**
No Past Breaks	11%	49%	1%	1%
One Past Break	41%	24%	98%	10%
Two Past Breaks	39%	12%	87%	4%
Three Past Breaks	9%	6%	87%	0%
Total Probability of Future Peace (Sum of Rows)				15%

I will not attempt a more balanced estimation approach here. However, we can make a very crude effort to correct for this lack of balance by shaving off *a quarter* of the probability drop. This last judgment call leaves us with approximately a **30% chance** of a peaceful future.

Current Outside View Probability Estimate: 30%

3.7.2 Data Uncertainty

We have already seen that all available long-run conflict datasets — including the Correlates of War dataset — have significant issues. We can think of them as including incomplete and likely biased samples of conflicts, drawn from a larger universe of conflicts. The Correlates of War dataset, in particular, essentially starts out as a dataset of conflict between European states and then gradually expands to include other regions. It also only extends backward to 1816. An "ideal dataset" would include all qualifying cases, for a reasonable definition of what constitutes an "interstate conflict."

Unfortunately, we will never have access to an ideal dataset. Still, though, it can be worth reflecting on what we would expect to find if we somehow *could* analyze an ideal dataset. Specifically, we should ask: How confident are we that our findings would be the

same, if we actually knew the "true" rate of interstate conflict for every year? If we are not confident that our findings would be the same, then our confidence in these findings should decrease.

The second appendix of this chapter attempts to clarify how accounting for data uncertainty should influence our overall forecast. *The main takeaway is that we are seemingly fairly safe to disregard data uncertainty.* For example, even if the rate of conflict in the nineteenth century was much greater than our dataset suggests, it is not clear that this would imply a large update to our forecast. Similarly, even if the world was much more violent before the start of the dataset, this would not clearly imply a large update either.

The basic explanation for these findings is the fact that most of the probability mass we currently assign to peace is driven by the hypothesis that *there has been a sharp downward trend in the rate of war since the middle of the twentieth century.* If this downward trend in fact dates back to the middle of the twentieth century, as our current dataset suggests, then data from long before this period has low relevance. This implies that our dataset's exclusion of many non-European states in the nineteenth and early twentieth centuries would not actually have much practical relevance. Similarly, the fact that our dataset does not cover the years before 1816 would not be a large problem.

On the other hand, suppose that our ideal dataset would reveal that the downward trend is in fact much older than it seems. In this case, we actually should become somewhat more optimistic. The less frequently trend breaks occur, the longer we should expect the current positive trend to last into the future. However, even if we suspect that the trend may be older than our dataset suggests, analysis in Appendix 2 suggests that fairly extreme assumptions will be needed to justify a *large* shift in our overall probability estimate.

Ultimately, data uncertainty does not seem to warrant substantially shifting our probability estimate from its current level.

Current Outside View Probability Estimate: 30%

3.7.3 Model Uncertainty

So far, we have been assuming that — except for occasional structural breaks, where our model's parameters change — our modified geometric Brownian motion model accurately describes the manner in which the rate of war evolves over time. However, the model is inevitably far from perfectly accurate. It is a three-parameter simplification of an enormously complex and opaque social phenomenon. We should expect its predictions to diverge from reality even if its parameters have been fitted as well as possible.

One standard way to account for model uncertainty is to consider several models, note what each of them predicts, and produce an overall prediction that is a weighted average of the predictions of the individual models. This is essentially the approach we have already taken, in considering variants of our model that display different numbers of trend breaks.

The first and most basic additional model we should consider is one *with the trend term removed*. This is, in effect, a simple way to give special weight to the null hypothesis that there is simply no trend.[37] In fact, calculating the Bayesian Information Criterion (BIC) for the trendless model, we find a value of 300.4. This is actually lower than the BIC score (302.4) for the single-trend-break model we discussed above. (BIC scores, recall, penalized more complex models.) Once we enter the trendless model into consideration, our earlier table of relative model likelihoods transforms into Table 3.7.

Crucially, about half of the probability mass is now dedicated to the version of the model that inherently lacks any trend. As a first pass response to this new finding, it may be

[37] One reason to give special weight to the null hypothesis is that it is generally reasonable — in line with Occam's razor — to have a bias toward giving simpler hypotheses and models additional weight. This is one intuitive justification for using standards such as the Bayesian Information Criterion to discriminate between models.

More speculatively: I think there may also be another distinct reason to want to downweight the models we have been using, beyond what values such as BIC scores alone would suggest. Essentially, it is very *easy* for the models we have been using to predict radical changes. Small differences in parameter values blow up to imply dramatic long-run changes in the world, without frictions or feedback loops to push against these changes. In practice, though, radical changes in the world seem to be very uncommon, because these kinds of frictions and feedback loops tend to exist in most domains. Giving some additional weight to a stable model can also be, in part, a crude way to give some additional weight to the broader hypothesis that the world may change less dramatically than toy models suggest.

Table 3.7: Results of Bai-Perron test, assessing the relative likelihoods of different models — including a model that lacks time trend term.

Model	**BIC score**	**Relative Probability**
Trendless Model	300.4	52%
Core Model, No Past Breaks	305.1	5%
Core Model, One Past Break	302.4	19%
Core Model, Two Past Breaks	302.5	18%
Core Model, Three Past Breaks	305.5	4%
Core Model, Four Past Breaks	320	0%

reasonable to remove approximately half of the probability we currently assign to a peaceful future, moving our overall estimate down from 30% to 15%.

However, recall, translating BIC scores directly into probability estimates rests on the assumption that we have no additional evidence or justifications for considering one model more plausible than another. Arguably, we do have some additional reasons to give weight to the hypothesis that the underlying tendency toward conflict has not been entirely unchanging over time.[38] Overall, then, we may not want to completely cut our probability estimate *entirely* in half. A new overall probability of estimate of something like 20% may be more reasonable.

Beyond considering model variants with *fewer* parameters, we can also consider model variants with *additional* parameters. Some options include incorporating the rates of conflict from years before the most recent year or adding polynomial terms that allow for non-linear effects. However, I will not actually explore any additional variants here.

More radical revisions of the model could also be considered. Options in this direction include changing the random noise in the model away from the normal distribution, adding an additional layer of random noise on top of the current model (e.g. using the output of

[38]Focusing specifically on the trend break models: the main period they appear as a breakpoint (the Second World War) would intuitively be a very unsurprising period for a break to occur. Stepping back, we also noted in Section 3.6.3.1 that qualitative histories of inter-European wars at least weakly suggest that the average rate of conflict may have at some point changed from its pre-nineteenth century value — given that both the "long peace" in the nineteenth century and the peaceful back-half of the twentieth century seem like they may genuinely have been historically unprecedented. These historical accounts, therefore, offer at least *weak* additional evidence that the expected rate of conflict is not entirely stable.

the current model as an input to a Poisson-distributed random variable), or moving away from models in the geometric Brownian motion family altogether.

Attempting to account for model uncertainty often has the effect of *increasing the uncertainty bars* on forecasts. Outcomes that initially seem almost guaranteed or almost impossible, according to a single model, can become less certain once we drop the assumption that the model is accurate. Since our current model suggests a *moderate* probability of a peaceful future, though, it is difficult to say whether considering additional models is more likely to increase or decrease our probability estimate.

Ultimately, to account for the fact that we should give the "no trend" model additional weight, I will adjust our estimate down from 30% to 20%. Although considering additional models could lead to further adjustments, I will not consider these additional models or make these further adjustments for now.

Final Outside View Probability Estimate: 20%

3.8 Summary: Where Our Forecast Stands

3.8.1 How to Interpret Our Outside View Forecast

We have now arrived at an overall "outside view" forecast of how likely the future is to be peaceful. Essentially, in producing this forecast, we have offered an answer to the following question:

> Suppose that we had virtually no insight into the causal mechanisms that determine the rate of conflict. Suppose, also, that we had virtually no information about the world other than information pertaining to the historical rate of conflict. *Under these restrictions*, what probability would it make sense to assign to the future being peaceful?[39]

[39]Of course, this is something of an exaggeration. In producing our forecast, we have not been able to

An outside view forecast should not be a stopping point. We do have *some* understanding of the causes of war, and we do have *some* further relevant information about the world. If we make a cautious effort to incorporate this understanding and this information into our forecast, we should — in principle — be able to do better. The next section of the paper will briefly comment on how we might go about this.

However, an outside view forecast can provide a useful foundation for further work. Studies of forecasts of social and political phenomena suggest that even *purely* "outside view" forecasts, such as simple trend extrapolations, often outperform forecasts that are not grounded in this kind of outside view (Tetlock, 2016). Successful forecasters often seem to work within a common framework: they *start out* by adopting the outside view, then move forward from there.

The specific number that we have arrived at is **20%**. This specific number is highly debatable, since it rests on a series of judgment calls. (Among these judgment calls, the most consequential and debatable judgment call was the way we chose to account for the expected impact of future trend breaks on the risk of conflict.) Even so, the number is not arbitrary. It is easy to argue that an alternative number would also be reasonable, but it is — I think — difficult to argue that a *dramatically* higher or lower number (e.g. 2% or 50%) would be reasonable; strong assumptions would be necessary to arrive at these kinds of numbers. We have at least made some progress in clarifying the spectrum of reasonable disagreement.

---

completely avoid thinking about causal mechanism or reflecting on information we have about the world. For example, in order to decide what metric to extrapolate forward, we needed to reason about which metrics were especially likely to display steady trends into the future; this required reasoning about causal relationships, for instance reasoning about whether changes in the number of states could partly account for past changes in the number of conflicts per year. The statistical model we fitted to the data was also at least partly theoretically motivated; we did not arrive at it simply by exhaustively testing possible models and comparing their fits. Nonetheless, so far we have certainly made *fewer* and *weaker* causal assumptions — and drawn on *less* information about the world — than one typically would when thinking about the future of conflict. In practice, the distinction between "outside view" and "inside view" analysis is more of a continuum than a binary.

3.8.2 How Our Choice of Question Has Mattered

Before moving on, it is worth pausing for a moment to remember our probability estimate concerns a very *specific* forecasting question:

> **Peaceful Future Hypothesis:** Between 2050 and 2030, the *average* rate of interstate conflicts involving the reciprocated use of force (per politically relevant dyad) will be below .0024.

If we had chosen a slightly different question, our estimate could be very different. For example, our estimate would presumably be higher if we had chosen a more generous threshold for the rate of conflict. (It is currently set at 10% of the apparent average rate between 1850 and 1900.) It would also be different if we chose a different window of time.

In general, the specific operationalization we have chosen does not necessarily perfectly correspond to what we care about. For example, I have chosen to focus on a very long and distant period of time (2050 to 2300) in part because I am interested in long-run trends in conflict and some of the more extreme possibilities that past writers have raised. If I had instead chosen to focus on the next thirty years, for example, then — seemingly — this period could easily be peaceful purely by chance. Choosing a longer window was necessary for the particular intellectual purpose of this paper. However, in practice, a forecast of the next thirty years is much more practically significant than a forecast of far-off decades.

Most importantly, in this vein, the future rate of "interstate conflicts involving the reciprocal use of force" could be very different from the future rate of more severe categories of conflict. For example, it is possible for the risk of great power war to stay the same even if the rate of our broader category of conflicts declines. On the other hand, it is possible for the risk of great power war to decline (e.g. due to the deterrent effect of nuclear weapons) even if there is no obvious decline in the broader category.

Regarding great power war, the unfortunate reality is that we simply do not have the statistical power necessary to detect trends in the rates of such rare forms of conflict. We

can gain a little bit of evidence from the patterns of spacings between rare events (e.g. how many years have occurred since the last great power war), but we cannot gain very much (Braumoeller, 2019).

As noted above, we can hope that the variable we are forecasting (the rate of interstate conflicts with reciprocated force) is actually closely intertwined with the variable we most care about (e.g. the rate of great power war). In other words, as expressed in Section 3.6.2.1.1, we can hope that the Constant Relative Probabilities assumption holds. However, it is not at all clear this is true.

A valuable exercise could be to examine the Constant Relative Probabilities assumption more closely, estimate how likely it is to be approximately correct, and then draw on this estimate to specifically forecast the future rate of (e.g.) great power war. Braumoeller's (2019) effort to identify changes in the distribution of casualty numbers is already one important step in this direction. However, this exercise lies beyond the scope of the current paper. (Appendix A will briefly outline some initial steps.)

3.9 Next Steps: Incorporating "Inside Views"

The outside view forecast tells us what we should predict, if we were to base our forecast entirely on trends in the rate of war. However, of course, we do not need to rely entirely on trends. We can also apply our understanding of the *causes* of war to shift our forecast up or down.[40]

As an analogy, suppose we were trying to forecast a business's revenue in the coming year. We could form an initial outside view forecast by extrapolating trends in its revenue.

[40] As an analogy, suppose we were trying to forecast a country's economic output next year. As a first past, we could generate an outside view forecast by simplifying extrapolating growth trends forward. However, we may not want to stop here. We would also want to consider changes in the world that could cause next year's output to exceed or undershoot the trend. As an extreme example, if we believed a country was about to experience an economically damaging disaster (e.g. a pandemic or an invasion), we would want to adjust our estimates downward; if the anticipated disaster was extreme enough, we might even want to almost totally disregard the existing trend. Other more mundane changes (e.g. change in a country's tax policy) might also justify more modest shifts.

We may want to adjust this forecast, though, by reasoning about factors that could cause the business's revenue to rise above or fall below the trend. For example, if we know that another major company will release a competing project in the coming year, then we can reason that this increase in competition will probably have a negative effect on revenue. We can then adjust our initial outside view forecast down, on the basis of this causal analysis.

Furthermore, as we have seen, it is actually highly ambiguous what the historical trends in conflict have actually been. We cannot *confidently* say there is actually currently a downward trend in the rate of conflict, or, if there is such a trend, when the trend actually started or how steep it actually is. Analyzing causal dynamics can also help influence our beliefs about the trends that already exist or have existed in the past.

As another analogy, suppose that some dataset suggests a recent sharp downward trend in a country's crime rate. It is unclear whether this is a real trend or whether it merely reflects a change in record-keeping or random noise. If we cannot identify any plausible causal explanation for the apparent downward trend, then we should assign less weight to the hypothesis that it is real.

Broadly speaking, we can take these same approaches to adjusting our forecast of the future rate of conflict. We can reflect on a variety of causal factors — ranging from the level of economic growth, to the prevalence of democracy, to the polarity of the international system, to the political aims of specific governments, to aspects of human psychology, and so on — and consider whether they give us any reason to adjust our forecast upward or down.

One way to approach the problem is to decompose it into a number of individual "causal stories" that need to be evaluated. These stories offer causal explanations for why the future may be likely or unlikely to be peaceful. For example, one story put forward by Angell (1913) is the "Trade Story," which suggests that rising economic entanglement could support a major decline in conflict.

The sub-problem of evaluating an individual "causal story" can then be decomposed into a three-step process. First, we can assess the relevant causal relationship (e.g. estimate

the effect of economic entanglement on the rate of conflict). Second, we forecast how the relevant variables will evolve (e.g. forecast how economically entangled states will become). Third, we can use these components to generate a *conditional forecast* about how the rate of war *would* evolve if there were no other relevant changes in the world.

Once we have evaluated several individual causal stories, and produced several individual conditional forecasts, we can then weigh them up to form an overall impression of the causal plausibility of a downward shift in war. This overall impression may then lead us to adjust our outside view forecast upward or down, if we are confident enough in this impression and it diverges strongly enough from what the outside view forecast suggests.

I will not pursue this line of within this paper. However, the path forward is *relatively* clear. Many of the necessary steps — such as identifying relevant "causal stories" and summarizing the body of the evidence for particular causal effects — have already been completed by other scholars. For example, Pinker (2011) and Gat (2017) both make major contributions to this project. However, a great deal of work remains to be done, including both (1) cataloging and integrating existing pieces of analysis into a combined forecasting framework and (2) conducting additional pieces of analysis to fill in gaps.

3.10 Conclusion

The field of International Relations was, at one time, deeply preoccupied with the *Peaceful Future Hypothesis*: the hypothesis that the world will become far more peaceful than it has been historically. This question of whether the future will be peaceful continues to haunt the field. From time to time, scholars still give themselves leeway to comment on the question. The question also seems to serve as an implicit motivation for some lines of research, such as Pinker's (2011) and Braumoeller's (2019) research on historical trends in conflict. However, as a topic of direct and sustained study, it is at best ambiguously respectable.

This paper has argued that the question is worth studying. Although we will never arrive at a definitive answer — or even a definitive probability estimate — we can still make meaningful intellectual progress. In particular, we can work to clarify and narrow the *spectrum of reasonable disagreement* about the probability of a peaceful future.

This paper has tried to clarify the Peaceful Future Hypothesis, clarify how it can be studied, clarify some of what we already know, and clarify what remains to be done. Along the way, the paper has offered an intermediate "outside view" forecast based on the analysis of historical trends. The specific number I arrive at is **20%**. While meaningfully higher or lower numbers would also be defensible, *dramatically* higher or lower numbers currently strike me as hard to defend.

To produce a more meaningful overall forecast, we still need to update our "outside view" forecast by also taking causal considerations into account. However, this next step is left for later work, including the subsequent chapters of this thesis.

3.11 Appendices

3.11.1 Appendix 1: Other Datasets

3.11.1.1 Expanded War Data

Gleditsch and Ward (1999; 2004) attempt to widen the geographic scope of the Correlates of War Interstate War dataset, by developing an expanded list of states and then cataloging wars involving these states. They aim to include any polity that:

> a) has a relatively autonomous administration over some territory, b) is considered a distinct entity by local actors or the state it is dependent on, and c) has a population greater than 250,000.

While their list of states largely converges with the COW list by the mid-twentieth century, it includes roughly twice as many states in the early nineteenth century. Some

examples of states it includes in 1816, which the COW dataset does not, are Japan, Haiti, Thailand, China, and Algeria. As a result, it includes a number of wars excluded by COW. These include the Opium Wars (fought between the European powers and China) and the Peruvian-Bolivian War of 1841.

For our purpose, however, there is an important tradeoff in using this data. Unlike the Correlates of War Militarized Interstate Disputes dataset, it does not include information about lower-intensity conflicts between the additional states. This means that we cannot, for example, use the dataset to track trends in the "reciprocated use of force" (Braumoeller's preferred metric). Using war datapoints alone will reduce our sample of cases and, all else equal, make trends more difficult to detect.

Arguably, although its coverage is more consistent and complete, the dataset still excludes a substantial number of polities that should be considered states. Griffiths and Butcher (2013) present a more up-to-date effort to catalog the world's states since 1816. They use the following standard as their criteria:

> A state is a recognizable political entity that possesses: 1. A population of at least 100,000. 2. Autonomy over a specific territory. 3. Sovereignty that is either uncontested or acknowledged by the relevant international actors.

Using these criteria, Griffiths and Butcher identify one hundred states not included in other datasets. For 1816, their list of states is *three times larger* than the Gleditsch and Ward list.

Examples of additional states that they include in 1816 are the Ashanti Empire, Zanzibar, Jaipur, and Bhutan. It also includes many states that are largely forgotten today, such as the central African Kuba Kingdom, which thrived for two centuries and was only reached by Europeans in 1884. A disproportionate number of the additions are for now-extinct African and South Asian states. In some cases, the additions reflect the slightly lower population threshold used by Griffiths and Butcher. However, in many cases, the states appear to meet the Gleditsch and Ward coding criteria as well.

Gleditsch and Ward's catalog of wars also appears to miss some nineteenth century wars fought between the states that it does recognize. For example, although it includes both Vietnam and Thailand as states since 1816, it does record the Siamese-Vietnamese War of 1831, in which Thailand attempted to conquer Southern Vietnam. In general, nearly all of the new wars it adds are wars involving Western powers.

Overall, the dataset therefore still appears to suffer from a weakened version of the same coverage issues as the COW dataset.

3.11.1.2 Correlates of War Non-State and Extra-State Wars Lists

The latest version (v4.0) of the Correlates of War Dataset has added a catalog of "non-state wars" (wars between non-state actors), in addition to "extra-state wars" (wars between states and non-state actors). Many of the "non-state actors" involved in these wars would, on a typical definition, be classified as states. For example, as noted above, China is classified as a non-state actor until the late nineteenth century.

Therefore, these catalogs essentially (at least in part) serve to expand our list of post-1816 interstate wars. We could in principle go through it and pick out the extra-state wars or non-state wars that might be more naturally classified as interstate wars. A key of limitation of this list, however, is that it does not come with a list of actors at any given time. This means that we do not know what the denominator should be, if we want to calculate the rate of interstate conflict. Another limitation is that, as in the case of the Expanded War Dataset, it does not include conflicts that fall short of war; this will likely make it substantially more difficult to analyze trends. A third limitation is the authors of the dataset simply do not claim that the catalog is comprehensive.

3.11.1.3 Conflict Catalog

Peter Brecke's Conflict Catalog attempts to catalog all "violent conflicts" throughout the world since 1400. Setting a fatality threshold of *at least thirty-two deaths in a single*

year, he defines "violent conflict" in the following way:

> An occurrence of purposive and lethal violence among two or more social groups pursuing conflicting political goals that results in fatalities, with at least one belligerent group organized under the command of authoritative leadership.

Brecke generated the dataset by drawing on "academic research manuscripts to encyclopedias by military historians to historical atlases to historical chronologies," including "major Chinese, Japanese, and Russian compilations that are essentially equivalent to what has been produced by military historians in the West except that they include many violent conflicts overlooked by Europeans and North Americans." Beyond extending back before 1816, the dataset contains a number of nineteenth century and early twentieth century conflicts excluded by both Correlates of War Interstate War Dataset and the Expanded War Dataset. For example, unlike Gleditsch and Ward's dataset, it does include the Siamese-Vietnamese War of 1831.

Unfortunately, for the purposes of extracting trends in interstate war, there are also a number of difficulties involved in using this dataset. One basic problem is that it does not disaggregate interstate wars and other kinds of violent conflicts. Interstate wars, civil wars, massacres, slave rebellions, and other diverse forms of violence are all mixed together.

It would be possible to "fix" this aggregation issue through additional coding work, to extract the subset of violent conflicts that qualify as interstate conflicts. However, even if we do this, there would still be a couple of barriers to calculating the rate of interstate conflict. There is both a "numerator problem" and a "denominator problem" here.

The numerator problem is that the dataset represents interstate conflicts as single events, which often pool together many dyadic conflicts. For example, the "Second World War in Europe" and the Falklands War are both simply listed as single events. Our preferred method of quantifying the rate of war, however, involves breaking down the Second World War into multiple dyadic components, instead of "counting it once." The dataset also does not seem to follow clear or consistent principles in pooling conflicts. Conceivably, though,

further disaggregation work – involving the consultation of the relevant sources – could mostly resolve this problem.

The "denominator problem," unfortunately, will be more difficult to address. The problem here is that the dataset does not include a list of the states that existed at each point in time. This means that we simply cannot calculate the rate of "war per politically relevant dyad," in a given year, because we do not know how many politically relevant dyads there were. The most comprehensive global dataset of states only extends back to 1816. Calculating the proper denominator, for wars extending back before 1816, would likely involve creating an entirely new dataset.

Putting these issues aside, the dataset's coverage undoubtedly becomes increasingly incomplete the further back in time it extends and the further it strays from Europe. If one small and long-extinct South Asian state killed thirty-two people in another small and long-extinct South Asian state, in 1420, there is no reason to assume that primary sources documenting this interaction would necessarily survive. Even when documentation does survive, there is also no reason to assume it would be reported and easily discoverable in the set of secondary source literature used to produce the Catalog.

A brief exploratory investigation confirms that even the Catalog's nineteenth century coverage is still fairly incomplete outside of Europe. To continue to use Vietnam as a test case: Its conflict with Champa – which previously occupied the territory that now makes up Southern Vietnam – are not included in the dataset, despite having stretched over hundreds of years and having only concluded in 1832. A telling general indicator of incompleteness is that the majority of states that existed in 1816 – according to the Griffiths and Butcher dataset – do not make any appearances in the Conflict Catalog, despite the Catalog having a very low official threshold for conflict inclusion. It is therefore plausible that the catalog still misses the substantial majority of conflicts even in the nineteenth century.

3.11.1.4 Historical Conflict Event Dataset

The Historical Conflict Event Dataset is another, more recent data that aims to cover conflicts – including non-European conflicts – for centuries before 1816 (Miller and Bakar, 2023).

It differs from the Conflict Catalog in two main ways. First, it focuses on "conflict events'' (e.g. battles) rather than conflicts. Second, it extends much further back in time: the earliest datapoint it includes is the Battle of Megiddo in 1468BC.

The Historical Conflict Event Dataset is largely (although not entirely) based on the set of conflicts included in Tony Jaques' multi-volume *Dictionary of Battles and Sieges: A Guide to 8,500 Battles from Antiquity through the Twenty-first Century* (Jacques, 2006). Jacques' dictionary in turn is largely based on the set of battles and sieges discussed in the (lengthy) Cambridge Modern History and Cambridge Ancient History book series, as well as Cambridge histories of individual countries.[41] It defines a battle as "any clash between organized forces of combatants," without setting any specific fatality threshold.

The dataset does include a number of battles from conflicts that are missed in the Conflict Catalog. However, even in relatively modern times, its coverage is clearly incomplete – and becomes more incomplete the further one moves backward in time and away from Europe. For instance, still focusing on Vietnam, the dataset records no battles involving Vietnam for the entire three-hundred year period from 1471 and 1773. In slightly more modern times, it also records no battles involving Vietnam between 1802 and 1847 – although Vietnam in fact experienced multiple wars and internal rebellions during this period (Kiernan, 2017). Meanwhile, as a reflection of its density grows as we approach the modern era and begin to include Western states, it records twenty battles for the period from 1964-1966 alone.

Ultimately, for our purposes, the Historical Conflict Event Dataset suffers from

[41] It ultimately includes the battles if they are cross-reference in at least two independent sources "with good detail or at least consensus among authorities on the date, the event, the participants and the outcome"

roughly the same issues as the Conflict Catalog: (a) it does not separate interstate conflicts out from other conflicts, (b) it does not give us all of the information we would need to calculate the rate of interstate conflicts per politically relevant dyad, and (c) its coverage is highly incomplete and becomes more incomplete the further one moves backward in time and away from Europe.

3.11.1.5 Further Datasets

Other conflict datasets that either extend further back in time or have more consistent and complete geographic coverage, relative to the COW dataset, have been compiled by Levy (1982), Clodfelter (2017), Wimmer and Min (2009), and Lyall (2020). However, for the sake of space, I will not discuss these alternative datasets and their merits and limitations here.

There are also alternative datasets whose coverage is more limited than COW, such as the UCDP/PRIO Armed Conflict Dataset that extends back to 1946 (Davies et al., 2023). Although these datasets may improve upon aspects of the Correlates of War dataset, they are too limited in temporal scope to support the identification of longrun trends.

3.11.2 Appendix 2: Accounting for Data Uncertainty

3.11.2.1 What if Our Dataset Included Non-European States From the Start?

As discussed above, one key problem with the Correlates of War dataset is that it largely begins as a dataset of inter-European conflict – excluding the majority of the world's states – and then gradually expands to become a global dataset of interstate conflict. This may bias our parameter estimates or create illusory trend breaks.

Histories of the nineteenth century generally suggest that the average rate of conflict was much higher *outside* of Europe in the middle and late nineteenth century than it was *inside* of Europe. For example, in Osterhammel's (2014) global history of the nineteenth century, the title of one chapter is simply: "Peaceful Europe, Wartorn Asia and Africa."

This raises the possibility of a longer-run global trend toward peace than either dataset

suggests. If we update our beliefs in the direction of believing that there *actually is* a single long-running downward trend, then this should make us more confident that the future will be peaceful (because it should make us more confident that the current downward trend will not break for a long time).

As a simple way to analyze our uncertainty, let us suppose that — if we could somehow rerun our statistical analysis on a complete dataset — we would either (i) continue to assign the same probability (i.e. 11%) to the "no trend break hypothesis" or (ii) instead come to assign it a 50% probability.

Given this simplifying assumption, we can create a new expression for $Pr(PeacefulFuture)$ that incorporates the possibility that a complete dataset would more strongly strongly suggests a single long-run trend without any past break:

$$\begin{aligned}&\Pr(\text{Peaceful Future}) = \\ &\quad \Pr(\text{Full Data Would Suggest } \textit{No Break}) \cdot \Pr(\text{Peaceful Future} \mid \text{Full Data Would Suggest } \textit{No Break}) \\ &+ \Pr(\text{Full Data Would Still Suggest } \textit{Break}) \cdot \Pr(\text{Peaceful Future} \mid \text{Full Data Would Still Suggest } \textit{Break})\end{aligned} \tag{3.7}$$

Although this estimate is extremely contestable, I think it could be reasonable to assign something like a 20% chance to the complete data suggesting a single trend rather than multiple trends with breaks. It would be easier for the model to accommodate the full sweep of the data — without a trend break — if the nineteenth century looked less peaceful. And it seems quite plausible that a complete dataset would *in fact* show the nineteenth century as being much less peaceful than it currently appears. However, what exactly the pattern would look like is very speculative.

Now that we have estimated Pr(Full Data Would Suggest *No Break*), we also need to estimate Pr(Peaceful Future|Full Data Would Suggest *No Break*). The way to do this is to estimate the *expected* probability that the single-trend model would assign to a peaceful future, if the data changes in a way that substantially improves the model's fit. We should

expect this probability to increase by some amount, but it is very hard to say by how much. One estimate might be that, in expectation, in this scenario, the single-trend model's new forecast would fall halfway between its current forecast (i.e. what the model forecasts when fitted to the Correlates of War dataset since 1816) and the current forecast of the one-break model (i.e. what the model forecasts when fitted to the Correlates of War dataset since 1937). This means that the model would assign a roughly 50% chance to a peaceful future.

Of course, these estimates are extremely arbitrary. However, it turns out, their arbitrariness does not matter very much. Plugging them into the relevant formulas implies that our overall forecast still approximately equals 30% (as it stood at the end of Section 3.7.1).

The upshot here is that issues with the Correlates of War dataset matter less than we might fear. Although there is no single obvious way to account for data uncertainty, it seems like we would need to make extreme assumptions for data uncertainty to substantially alter our overall probability estimate. Data uncertainty, seemingly, should not lead us to change our overall probability estimate by very much.

3.11.2.2 What if the Correlates of War Dataset Extended Further Back?

The other seemingly important limitation of the Correlates of War dataset is that it only extends back to 1816. We can also try to understand how our overall forecast might change if we had access to data that extended further back in time.

There are three basic ways this longer-run data could influence our overall forecast:

1. **Changing our long-run trend estimate:** First, if we make the assumption that the "no historical trend breaks" model is correct, then having access to this earlier data would influence our estimate of the current trend.
2. **Changing our estimates of the likelihood of future trend breaks:** Second, by observing the rate of trend breaks in the earlier data, we could arrive at better estimates of how long current trends will persist.

3. **Informing predictions about the characteristics of trend breaks:** Third, by observing additional trend breaks, we might be able to develop more informed predictions about the potential magnitudes and directions of future trend breaks. For example, if we find that trend breaks tend to result in radical changes in direction, then this would have forecasting implications.

Let's start with the first consideration. It seems likely that extending our dataset backward — at least for the (primarily European) region the dataset starts in — would lead to a downward trend across the dataset. We know that the thirty years that preceded the dataset's start were periods of intense conflict in Europe, corresponding to the French Revolutionary Wars and Napoleonic Wars. We also know that contemporary observers considered the period after the Napoleonic Wars to be anomalously peaceful in European history. However, *insofar as we are still assigning a relatively low probability to the existence of a single long-run trend*, this observation about a trend in a previous era should not have much direct influence our forecast.

Now, let's move onto the second consideration. As an illustration of how pre-1816 trend breaks can matter, we can consider how it would influence our forecast to find either zero or two clear trend breaks between 1700 and 1816. Figure 3.12 shows the new implied probability distributions for how long we should expect the current trends to last. We can see that different findings can meaningfully shift our persistence forecasts, especially if we make the assumption that there have been no breaks since 1816.

To what extent should we actually expect to find trend breaks in the pre-1816 data, though? I see no obvious reason to assume that there were either very many trend breaks in the past or very few. We could attempt to produce an estimate by assuming that major historical events (e.g. the French Revolution) tend to result in trend changes, but this seems like a dubious assumption. Given that we cannot easily identify trend breaks even by looking at graphs, it is hard to imagine that we could reliably identify trend breaks by considering qualitative historical accounts.

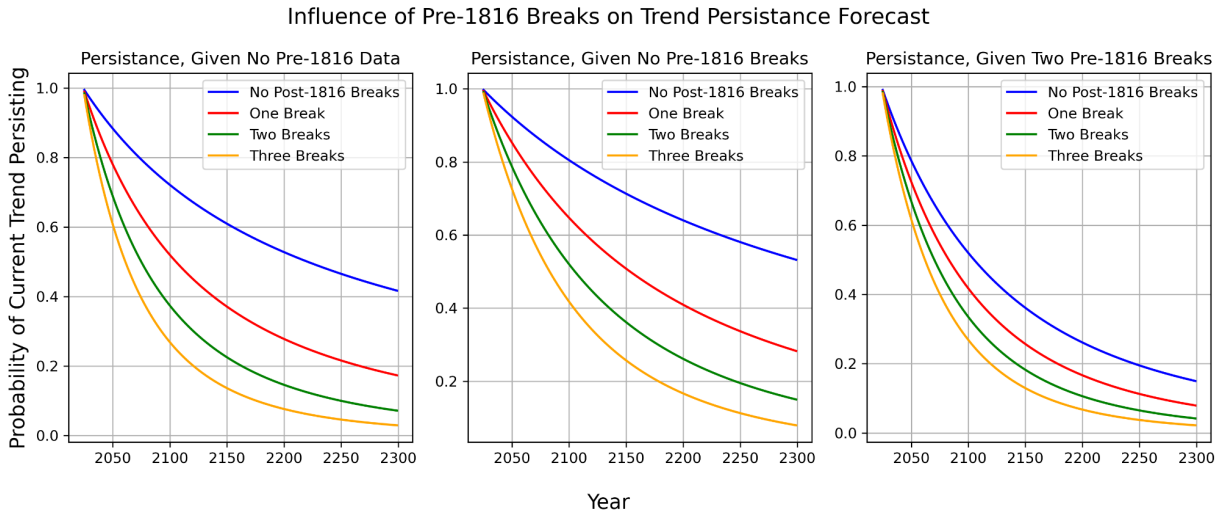

Figure 3.12: The likelihood of the current trend in conflict persisting, given different assumptions about past trend breaks. The first plot assumes we have no information about trend breaks in the period from 1700 to 1816, a period for which we largely lack data. The second plot is based on the assumption that there was a single trend break in this period. The third plot is based on the assumption that there were two trend breaks in this period. We see that assumptions about the more distant past can have some influence on our forecast.

The single thing I think we can say, on the basis of qualitative accounts, is that it is reasonably likely there was *at least one* trend break *sometime* in the centuries before the mid-nineteenth century. If we trust that the average rate of conflict in nineteenth century Europe was in fact extremely anomalously low, then that gives us at least some reason to think that was a genuine downward trend in the period leading up to it (and perhaps during it as well). This trend cannot have started too many centuries earlier, because, if it had, then the rate of conflict in the distant past would need to have been impossibly high. Insofar as we buy this line of argument, we should shift our most optimistic trend persistent estimates slightly downward and — as a result — become slightly less hopeful that the future will be peaceful. *However, this shift in our probability estimates should not be very large.* If we assume a single trend break at some point within an additional few-century period, this should only have a relatively modest impact on our persistence curves; the argument for the existence of this break is also fairly speculative.

Finally, let's now turn to the third consideration. So far, we have modeled the effects

of future trend breaks in an extremely crude way. We have been starting out with the asymmetric assumption that future trend breaks can only lead to a much larger risk of conflict, calculating how much of the value of Pr(Peaceful Future) would then be lost because of possible future trend breaks, and then — as a rough way of trying to correct for the asymmetry of our assumption — restoring a quarter of the lost probability. We are doing this because we have virtually no information about the historical patterns in the magnitudes or directions of trend breaks. If we had this information, though, then we could model the effects of future trend breaks in a more sophisticated way.

If the pre-1816 data does not contain many examples of trend breaks, then we will not learn much. However, if it contains many examples, then we could begin to estimate a distribution over the direction and magnitude of changes in our model parameters. For instance, if we find that trend changes have typically been *downwards* over the past several hundred years, then this would of course be at least some cause for greater optimism. If we find that trend breaks tend to have very large magnitudes and no obvious directional bias, then this would imply high variance in our estimates of the future rate of war. Ultimately, though, we certainly do not seem to be in a position to estimate the sizes and magnitudes of past trend breaks, given that we are not even in a position to identify these breaks. *This is, seemingly, another form of uncertainty that should not lead us to meaningfully alter our current forecast.*

To summarize, if we had access to longer-run data:

- We would probably find that the rate of war was meaningfully higher before 1816. However, this anticipated finding should not influence our forecast very much. The simple reason is that there is likely to have been a structural break between 1816 and today.

- We would also gain more information about the historical rate of structural breaks and the characteristics of these breaks. This information could be highly relevant to our forecasts. However, qualitative historical accounts do not provide us with a strong basis

for even speculating about past trend breaks. Therefore, reflecting on the possibility of these breaks should not lead us to update our forecast very much either. The one change that may be justified is a slight decrease in our overall probability estimate, based on arguments for at *least a single trend break* having occurred sometime before the nineteenth century.

In short, even if we attempt to correct for the fact that the Correlates of War dataset excludes older conflicts, we should not update our overall forecast very much. *More generally, we lose surprisingly little by simply disregarding data uncertainty.*

Chapter 4

Growth and Civil War

4.1 Abstract

Economic growth seems to reduce the risk of civil war, but it is not clear why. This paper tries to improve our understanding of the link between growth and peace, by exploring a possible causal mechanism that has mostly been neglected. I argue that, because there are strongly diminishing returns on consumption, wealthier people have less to gain from conflict. They assign increasingly little value to the additional income they can earn through violence, even if it grows in proportion to the wealth of their society. This shift in incentives has probably helped to lower the risk of civil war. We also should not be surprised if it has implications for interstate war.

4.2 Introduction

Wealthier states experience fewer civil wars. This relationship has been established statistically, but it is also large enough to be noticeable even without statistical analysis. When one of the world's most impoverished state falls into civil war, observers tend to react with sadness but not with shock. If Switzerland were to fall into civil war, on the other

hand, then there *would* be shock.[1]

Wealth and conflict are probably linked in multiple ways. For example, historical damage from colonialism may simultaneously fuel conflict and limit growth. Conflict — and the anticipation of conflict — can also clearly limit growth.

Still, when interpreting the statistical evidence, most scholars of civil war conclude that economic growth probably does have a pacifying effect. Growth seems to actually make civil war less likely. The main question is *why*. Scholars have discussed many possible causal pathways, including pathways that run through increases in state capacity and pathways that run through changes in physical infrastructure.

In this paper, I argue that one comparatively "direct" causal pathway has largely been neglected. In particular, economic growth may reduce the risk of civil war simply by making people wealthier and therefore more satisfied with their income levels. Since there are strongly diminishing returns on income, the potential spoils of war are less able to motivate wealthy people to risk their lives or their freedom.

To lay out this argument in slightly more detail: Prosperity seemingly has two opposing effects, which we should not expect to perfectly balance out. First, prosperity increases the potential spoils of conflict: the wealthier your neighbor is, the more you could win by robbing or extorting him. Second, prosperity makes people more satisfied with what they have: the wealthier you are, the less you value each additional dollar you could take from your neighbor. The relative strengths of these effects depend on just how quickly people become satiated. Available empirical evidence — which will be discussed below — suggests that people do in fact become satiated quickly enough for conflict to lose appeal.

The first section of the paper briefly characterizes the country-level statistical association between prosperity and peace. It then briefly summarizes the main causal mechanisms that scholars have proposed to explain the association. The *Diminishing Returns Mechanism* I have just described has been raised before in the literature, but has generally been

[1]For reference, for the year 2007 — the most recent year covered by the Correlates of War dataset — the set of states experiencing civil wars was Colombia, Somalia, Yemen, and Sri Lanka.

dismissed. For instance, a foundational paper in the civil war literature — Fearon (2008) — briefly raises and then rejects the idea that diminishing returns are relevant to insurgents' financial incentives.

The second section of the paper clarifies and provides support for the Diminishing Returns Mechanism. It uses rational actor models to explore the relationship between prosperity and individuals' incentives to participate in risky conflicts. It shows that the net effect of prosperity depends on how quickly the returns on consumption diminish. All else equal, evenly distributed economic growth will have an aggravating effect if the returns on consumption are *more-than-logarithmic* and will have a pacifying effect if the returns on consumption are *less-than-logarithmic* (even if they are nonetheless diminishing to some degree). Fortunately, empirical evidence suggests that they are sublogarithmic. Taking non-financial motives for war into account complicates the picture, but does not reverse the ultimate conclusion. Prosperity probably tends to reduce individual incentives to fight. Therefore, prosperity probably reduces the risk of civil war.

The third section discusses additional sources of empirical evidence that bear on the validity Diminishing Returns Mechanism, including studies of within-country variation in levels of violence. Overall, some of this additional evidence supports the hypothesis and some of it is in tension with the hypothesis. Although it is ambiguous how this evidence "nets out," I do not believe it should lead us to radically revise our views in either direction.

Ultimately: As in well-known, growth seems to reduce the risk of civil war. There are multiple well-known reasons why growth might have this effect. However, the Diminishing Returns Mechanism outlined in this paper could plausibly be a substantial part of the story.

4.3 Background

4.3.1 Prosperity and Peace: The Country-Level Statistical Relationship

There is a strong and robust negative association between GDP-per-capita and the risk of civil war. In this section, I will briefly characterize this association.

In the civil war literature, an especially foundational and widely-acknowledged finding is: *Wealthier states are much less likely to experience civil wars* (Fearon 2003; Hegre 2006; Ray 2017). The relationship is robust to the inclusion of controls that represent a number of possible confounders, such as the extent to which the country has a democratic government.

While the relationship can be studied using a number of different research designs — with different standards for what counts as a "civil war," different control variables, different sources of data, and so on — Table 4.1 presents representative results from a set of logistic regressions on the Correlates of War intrastate war dataset.[2] The key independent variable here is Log(GDPPC), the logarithm of the country's GDP-per-capita. The results of the regressions are consistent with the academic consensus: wealth predicts intrastate peace. This finding holds whether we look all the way back to the early nineteenth century or whether we simply look at the post-WWII period.

Of course, however, the correlation here could be spurious. One concern is that states may have certain fixed traits that help to determine both how wealthy they are and how peaceful they are. Economists have identified dozens of stable or slow-to-change geographic, institutional, and cultural factors that are predictive of a country's level of development (Weil, 2009). Examples of such factors include whether a country is landlocked, whether a country has large natural resource endowments, whether a country has a high-trust culture,

[2]The dataframe for these regressions was constructed using the peacesciencer package in R, which combines the Correlates of War dataset with a number of other data sources (Miller, 2022). The reader can refer to the package documentation for all sources. The GDP-per-capita data, for example, is derived from Anders et al. (2020), who in turn construct the early years of their dataset primarily by drawing on a previous dataset by Maddison (2010).

Table 4.1: Logistic regression model without fixed effects.

	Dependent variable:			
	Civil War Outbreak			
	(1)	(2)	(3)	(4)
Log(GDPPC)	-0.168^{****}	-0.205^{****}	-0.426^{****}	-0.770^{****}
	(0.028)	(0.034)	(0.051)	(0.116)
Democracy		-0.036^{****}		0.017
		(0.011)		(0.016)
Log(Population)		0.351^{****}		0.293^{****}
		(0.039)		(0.066)
EthnicPolarization				0.931^{**}
				(0.448)
Log(PercentMountains)				0.158^{**}
				(0.076)
MinYear	1817	1817	1947	1947
Observations	13,888	12,787	8,580	7,186
Note:		*p<0.1; **p<0.05; ***p<0.01; ****p<0.001		

and whether a country has a history of colonialism.

Another concern is that there may be time trends in the rate of civil war that are unrelated – or only partly related – to our control variables. For example, during the Cold War, the bipolar structure of the international system may have had an influence on the rate of civil war initiation. Since the average wealth of countries has been rising over time, unrelated time trends in influential global variables (e.g. the polarity of the international system or the popularity of certain ideologies) could lead to spurious correlations between wealth and conflict.

We can partly address both of these concerns by using *two-way fixed effect* models. These models include dummy variables for each state, as well as dummy variables for each time period. They control for any fixed traits of a particular state that make it more or less

prone to civil war, as well as for any traits of a particular time period that makes civil wars more or less likely to break out during it.

This leads us to our second, stronger finding: *Within a given time period, states become much less likely to experience civil wars as they become wealthier.* In order words, wealth is still strongly negatively associated with conflict even once we include dummy variables for each state and for each year. Table 4.2 reports the results of a set of regressions with two-way fixed effects.[3]

As can be seen in Table 4.2, one caveat to the finding is that the relationship between growth and peace only remains statistically significant when we look over a long window over time. When we restrict our attention to the post-WWII era, it dips below the statistical significance threshold. However, this lack of statistical significance for the post-WWII sample may simply reflect a decline in sample size and therefore in statistical power. Since the confidence intervals for the full-sample estimate and the post-WWII estimate substantially overlap, our findings are still *consistent* with a relationship that has been fairly steady over time.[4]

Of course, even if we accept that states become less likely to experience civil wars as they grow, the finding does not imply that there is *necessarily* a causal relationship here. The key problem is that we are still not, for the most part, controlling for traits of individual states that change over time. For example, suppose that female empowerment both reduces the risk of civil war and contributes to economic growth. Female empowerment certainly changes over time within a given state (i.e. it is not a fixed trait of states) and certainly varies from state-to-state within a given year (i.e. it is not a fixed trait of years). Therefore, our dummy variables mostly do not control for it. We could, if we wanted, try to eliminate female empowerment as a possible confounder by introducing a metric for female empowerment as a

[3]These regressions were run using the feglm function from the R package fixest. Standard errors are clustered by country.

[4]Unfortunately, alternative civil war datasets (such as Uppsala Conflict Data Program dataset) tend to only cover the post-WWII era. Due to this limitation, they tend not to show a statistically significant relationship between growth and peace when fixed effects are used.

Table 4.2: Logistic regression model with dyad-level and year-level fixed effects.

	Dependent variable:			
	Civil War Outbreak			
	(1)	(2)	(3)	(4)
Log(GDPPC)	-0.941****	-0.946****	-0.497	-0.475
	(0.249)	(0.269)	(0.390)	(0.383)
Democracy		0.038**		0.042
		(0.018)		(0.026)
Log(Population)		0.139		-0.764
		(0.338)		(0.677)
EthnicPolarization				-1.717
				(3.211)
MinYear	1817	1817	1946	1946
Observations	7,245	7,136	2,864	2,680
Note:	*p<0.1; **p<0.05; ***p<0.01; ****p<0.001			

control variable. However, realistically, we cannot do this for every possible uncontrolled-for confounder; our strategy will also be limited by our ability to define good metrics and collect reliable data. Therefore, inevitably, there will still be room for doubt about whether wealth actually *causes* peace or whether the statistical association is spurious.[5][6][7]

[5]One particular practical issue here is that most existing political science datasets that we might want to draw on to introduce controls — for example, the ethnic polarization dataset we have used — only cover the post-WWII period, during which the relationship between growth and peace is already statistically significant.

[6]Djankov and Reynal-Querol (2010) come to the contrary conclusion that the use of fixed effects alone is actually sufficient to make the association no longer significant. They arrive at this result, however, in substantial part because they bucket their dataset into 25-year-intervals. (This means that there are only four datapoints for each state rather than approximately 200.) Using this research design, they fail to find a statistically significant relationship between wealth and peace. Replicating their analysis (using more up-to-date GDP data), I arrive at the same finding, but find that the relationship again becomes statistically significant as we shrink the interval lengths back down. An interval size of 10 years, for example, is sufficient for the relationship to be significant at the level of p = .03. Of course, as shown in Table 4.2, simply using a single year interval implies a highly significant relationship. Overall, then, this contrary finding may mainly reflect the fact that bucketing reduces statistical power.

[7]Some authors have also approached the country-level causal inference problem through the use of instrumental variable methods. A paper of particular note is Miguel et al. (2004). They use an instrumental variable model, with rainfall as the instrument, to search for a causal link between GDP-per-capia and the

Overall, even though there is still room for doubt, the statistical association between growth and peace does at least seem to *suggest* that growth reduces the risk of civil war. The association is unusually strong and unusually robust to the inclusion of controls, notably controls for state-level and year-level fixed effects.

4.3.2 Possible Mechanisms Linking Growth to Peace

4.3.2.1 The Main Proposed Mechanisms

When countries become wealthier, they also tend to undergo a wide range of *structural* economic and political changes. For example, populations tend to migrate from the countryside to the city, the agricultural sector tends to decline in importance, infrastructure tends to improve, education levels tend to rise, and so on. For the most part, when scholars of civil war have asked why growth may cause of peace, they have focused their attention on causal pathways that *run through* changes in these kinds of structural variables.

In particular, scholars of civil war have dedicated the bulk of their attention to causal pathways that run through (1) changes in state capacity, (2) changes in terrain, (3) changes in the character of wealth, and (4) changes in the relative prices of different commodities (Ray, 2017). They have noted, *first*, that development tends to increase the power of state actors relative to non-state actors, thus reducing the feasibility of insurgency (Fearon, 2008). They have also noted, *second*, that it leads to improvements in infrastructure and reductions in wilderness that further disadvantage insurgent groups. They have further noted, *third*, that development leads to a shift in the portion of a country's wealth that can easily be seized.[8] It may be relatively easy to seize a piece of land without destroying much of its

risk of conflict. The authors find that higher growth *rates* lead to lower rates of conflict, but fail to find a significant relationship between GDP-per-capita and rates of conflict once growth rates are controlled for. (Later commentaries, particularly Ciccone (2011), have also questioned the link between growth rates and peace.) Since Miguel et al.'s rainfall dataset only allows them to examine a short sliver of the post-WWII period (1981-1991), however, we would not necessarily expect to see a statistically significant relation. Furthermore, instrumental variables models that use rainfall as an instrument have since come under significant criticism (Mellon, 2022).

[8]Note, however, that growth should still tend to increase the *total* amount of wealth that can be seized.

value, or to seize items from another person's home without damaging them, but, as an extreme example, it is relatively difficult to seize the wealth represented by Facebook's stock value. *Fourth*, scholars have noted that economic growth may be uneven across sectors, in a way that also differentially affects the amount of money that potential combatants can earn peacefully and the amount of money that they can be earned through looting (Dal Bó 2011).

4.3.2.2 The Diminishing Returns Mechanism

In this paper, we are interested in the possibility of another, more direct mechanism linking wealth to peace. In particular, we are interested in whether wealth might drive peace by making people more satisfied with their current income levels and, therefore, less easily enticed by the potential spoils of war. We can call this possible mechanism the *Diminishing Returns Mechanism*, since it hinges on the extent to which people experience diminishing returns on income.

In a classic paper on the causes of civil war, Fearon (2008) does consider this mechanism. He raises the possibility that, due to diminishing returns, rising prosperity reduces the economic incentive to become an insurgent. However, he ultimately dismisses this possibility. He argues that diminishing returns could only reduce incentives for conflict if people exhibited an implausible tendency known as "increasing relative risk aversion." If we instead make the more plausible assumption that people exhibit "constant relative risk aversion," then, he argues, prosperity will tend to increase both the subjective costs and benefits of conflict equally.[9]

Most reviews of the literature on economic development seem to share this perspective that, regardless of whether people experience diminishing returns on consumption, the effects of evenly distributed growth should wash out (Dixon 2009; Blattman 2010; Ray 2017). For instance, in a survey article for *Annual Reviews*, Ray (2017) briefly notes that rising incomes are "obviously inadequate" to explain the observed negative correlation between a GDP-per-

[9]This objection will be discussed below.

capita and civil war, because evenly distributed income growth should increase both the subjective costs and benefits of conflict participation equally.

Gat (2017), focusing instead on the case of *interstate* war, is perhaps the most prominent scholar to have argued that diminishing returns are relevant to the risk of conflict. However, he does not formalize his argument or explore it in depth. For example, he does not note that the satiation effect must be strong enough to compensate for the fact that growth also increases the total amount of wealth that can be seized. He also does not explore the implications of different returns curves; as we will note below, it matters exactly how *sharply* the returns on consumption diminish.

Overall, the civil war literature has largely dismissed the possibility of the Diminishing Returns Mechanism. Political scientists have also never explored this potential mechanisms in significant depth.[10]

4.4 The Diminishing Returns Mechanism

Individuals can, of course, have a variety of incentives to participate in civil wars. For example, a combatant might hope to enrich themselves by collecting a salary, by looting, or by opening up new career opportunities. They might also hope to increase their social standing, right moral wrongs, escape unpleasant domestic situations, or experience excitement and camaraderie. Non-combatants can also have similar, although typically weaker, incentives to advocate for or otherwise lend their support to wars. On the other hand, wars can of course pose substantial risks to both combatants and non-combatants alike. Death and internment are two especially notable concerns.

[10]On the other hand, some economists have conducted in-depth analysis of the behavioral implications of prosperity in other domains. Hall (2007) has suggested that rising prosperity helps to explain rising healthcare spending as a portion of GDP. He argues that wealthy people should be willing to trade away increasingly large portions of their incomes for the sake of avoiding early impairment or death. Jones (2016) has also suggested that rising prosperity may lead to declining support for innovation, as people become increasingly satisfied with what they already have and increasingly unwilling to tolerate dangers posed by new technologies. Finally, Aschenbrenner (2019) has suggested that rising prosperity may result in a rise in spending to address global risks such as climate change and new pandemics. The theoretical framework in this paper is most strongly influenced by Jones' work on innovation.

In this section, using simple rational actor models, I explore the relationship between wealth and rational willingness to tolerate these risks. The most complete model described in this section incorporates both financial motives and non-financial motives such as personal honor.[11]

These models serve to formalize and clarify the Diminishing Returns Mechanism sketched above. Overall, they provide support for the possibility that the mechanism does play an important role in reducing the risk of civil war as countries grow.

4.4.1 Risk-Taking and Financial Incentives: The Core Model

4.4.1.1 The Value of Consumption

First, let us consider models in which an individual's utility is a function of their consumption.[12][13] The functional form that is most commonly used when studying the relationship between consumption and utility is the isoelastic utility function (Layard, 2008):

$$U(C) = A + \frac{C^{1-\gamma}}{1-\gamma} \tag{4.1}$$

Here, C is the individual's consumption level. Since life is not possible below the subsistence level of consumption, C_{sub}, it is assumed that $C \geq C_{sub}$. γ is a parameter that describes how quickly the returns on consumption diminish. When $\gamma = 0$, for instance, the

[11] These models assume that people can freely choose whether to participants in conflicts. Of course, this is not always true. In many conflicts, a substantial portion of participants have experienced abduction. Even in these cases, though, these kinds of financial or non-financial incentives do presumably tend to influence how acquiescent individuals are or how motivated they are to pursue opportunities to leave the groups they belong to.

[12] One simple definition of *consumption* is that it is a measure of personal spending on goods and services. We can also define it more abstractly: as a measure of how much one would need to spend, in a market economy, to buy the goods and services that one uses. This more abstract definition emphasizes that we are primarily interested in what an individual is using, rather than how the individual acquired these things. An ideal measure of consumption would, for example, be insensitive to whether someone bought a given article of clothing, received it as a gift, or sewed it themselves. Consumption is of course closely related to income: Although income can be invested rather than consumed, higher incomes ultimately enable higher levels of consumption.

[13] Although consumption is far from the only factor that influences utility, it is certainly an especially important one. Differences in average incomes account for the majority of cross-country variation in average life satisfaction (Helliwell, 2012).

utility function is linear and the individual experiences no diminishing returns. As $\gamma \to 1$, the utility function increasingly approximates a logarithmic utility function. For $\gamma > 1$, returns on consumption are sub-logarithmic.

A is a constant term. When $\gamma < 1$, the second term in the equation is positive and increasing in magnitude, and A serves as a loose lower bound on utility. When $\gamma > 1$, the second term in the equation is negative and decreasing in magnitude, and A instead serves as a tight upper bound on utility.[14]

While the value of A has no inherent implications for the decisions an individual makes, we can fix its value, without loss of generality, by associating the state of being dead with 0 utility (Jones 2016). Let us also assume that people prefer life at the subsistence level to death. This means that $U(C_{sub}) > 0$. The higher the value of A is, for a given value of γ, the more people value subsistence living relative to death.

The isoelastic utility function has a number of different points of appeal. First, the function is very flexible. The parameter γ can be adjusted to produce a wide variety of returns curves. Second, the function is simple and easy to analyze. Third, the function is often considered at least an acceptably good match for observed human behavior.

4.4.1.2 The Disvalue of Death and Prison

In the context of expected utility models, death and internment can be regarded as special cases of what we will call *leveling risks*: risks that threaten to decrease an individual's utility down to some fixed value L. In the case of death, for instance, the individual's utility drops down to zero. Imprisonment can be modeled as an event that decreases an individual's consumption level down to some low fixed level C_L, set by their captors, and thus decreases their utility to the fixed level $U(C_L)$.

Notably, the higher an individual's level of utility rises, the more utility a leveling risk would destroy. In this sense, well-off individuals tend to have "more to lose" from any

[14]One essential difference between these two cases is that, only when $\gamma > 1$, the returns on consumption decrease quickly enough to imply the existence of such an upper bound.

leveling risks posed by war.[15]

4.4.1.3 Conflict as a Gamble for Wealth

4.4.1.3.1 A Gamble With a Fixed Prize Now, given this set-up, we can analyze the behavior of individuals who face leveling risks such as internment or death. Consider someone who is faced with two options: a status quo option and a risky option. If they take the status quo option, then they maintain their current level of consumption C. If they take the risky option, then they have some probability $p < 1$ of increasing their consumption by ΔC. Unfortunately, if they take the risky option, there is also a probability $(1 - p)$ that their utility will decrease to some lower fixed level L.

They will take the risky option if its expected upside is larger than its expected downside:

$$p[U(C + \Delta C) - U(C)] > (1 - p)[U(C) - L] \tag{4.2}$$

How does the individual's initial consumption level affect their decision? The answer depends on the value of the risk attitude parameter γ.

For $\gamma \geq 0$, the individual experiences diminishing or linear returns on consumption. This means that, as consumption grows, the difference between $U(C+\Delta C)$ and $U(C)$ either shrinks or remains the same. On the other hand, as consumption grows, the difference between $U(C)$ and L grows too. The downside of the gamble rises, while the upside does not. It can be shown that, for $\gamma \geq 0$, there is inevitably some level of initial consumption above which the person will not take the gamble.

4.4.1.3.2 A Gamble With a Proportionate Prize We can also consider the case where the size of the "prize" is not fixed, but rather is proportional to the individual's initial

[15]This model implicitly assumes that, when people decide how willing they are to risk death, they at least implicitly take into account what they will be missing out on if they die. However, as we will discuss below, other plausible ways of modeling death ultimately yield similar results.

consumption level. For example, we can ask how the attractiveness of "double or death" gambles rises with consumption.

These kinds of cases are particularly relevant to decisions concerning participation in civil wars, since broadly distributed economic growth will also tend to increase the salaries that political actors can afford to pay combatants or the wealth that exists to be looted.[16] We can think of the prize here as, essentially, the wage premium a combatant will earn if they are not killed or locked up.

Let $R > 1$ be the factor by which the individual's consumption will increase if they win the prize. Then the relevant inequality becomes:

$$p[U(RC) - U(C)] > (1 - p)[U(C) - L] \tag{4.3}$$

If $\gamma > 1$, then there are sub-logarithmic returns on consumption, and the difference between $U(RC)$ and $U(C)$ shrinks as consumption grows. It can be shown that, for any value of R, there will be a level of consumption beyond which the expected upside is smaller than the expected downside.

If $\gamma < 1$, then returns diminish more slowly. The difference between $U(RC)$ and $U(C)$ actually becomes larger for larger values of C. The question, then, is whether the expected upside or the expected downside of the gamble grows more quickly. This will depend on just how large the values of p and R are.

When $pR^{1-\gamma} < 1$, the upside grows more slowly and the gamble eventually becomes irrational. When $pR^{1-\gamma} > 1$, the upside instead grows more quickly.[17]

[16] Economic growth within a country does tend be fairly broadly distributed and does not on average change the distribution of wealth very much. Dollar (2016) reports that, when a country experiences growth, "[a]verage incomes in the poorest two quintiles on average increase at the same rate as overall average incomes."

[17] When $pR^{1-\gamma} = 1$, both the upside and the downside grow at the same pace. Notably, the individual's consumption level is irrelevant to the value of the gamble only when this knife's edge condition is met.

4.4.1.4 Empirical Evidence

4.4.1.4.1 Parameter Estimates The parameter γ represents the speed with which the returns on consumption diminish. We have seen that it determines the relationship between prosperity and the incentives potential combatants have to tolerate "leveling risks" such as death and internment.

There are three main approaches to estimating the typical value of γ.[18]

First, researchers can observe the economic behavior of individuals to directly estimate γ. For example, they can analyze the level of risk aversion implied by the portfolios of investors. As Jones (2016) reports, this body of literature consistently suggests that $\gamma > 1$.

Second, under the assumption that self-reported life-satisfaction is a rough proxy for an individual's level of utility, researchers can observe the relationship between income and self-reported life-satisfaction (Layard, 2008). These estimates, which are also primarily cross-sectional in nature, typically suggest that $\gamma > 1$.

Third, researchers can attempt to estimate the income elasticity of the statistical value of life. The statistical value of life (VOL) is a measure of the financial compensation required to entice an individual to accept a marginal increase in the probability of death, divided by the size of the increase (Cropper, 2011). It is typically estimated by running regressions on datasets that include information about the salaries and mortality risks of various jobs; the positive relationship between salary and risk is assumed to reflect a risk premium, once other job characteristics are controlled for. The VOL is also sometimes estimated by asking individuals to judge hypothetical trade-offs between income and mortality risk. The income elasticity of the VOL is then typically estimated either by comparing VOL estimates for countries with different average income levels or by comparing VOL estimates within a single country as it grows wealthier. As Kaplow (2005) explains, under certain weak assumptions,

[18]There is also substantial evidence that people's responses to risk are sensitive to contextual factors that our simple rational actor model excludes (Levy-Garboua 2012). Even individual people do not act as though γ has some single, consistent value. Discussions of the value of γ are therefore best-interpreted as discussions of typical ranges of values. The common claim that $\gamma > 0$, for instance, is best-interpreted as meaning something like: "A typical person typically displays some degree of risk aversion."

the elasticity of the VOL serves as an upper bound on γ.

Most studies that cross-compare countries with different average levels of income find elasticities less than one. However, the two notable studies that use panel data for individual countries (the United States and Taiwan) find elasticities greater than one (Hammitt 2000; Costa 2004). This discrepancy suggests that the cross-sectional studies may suffer from omitted variable bias.

Ultimately, Jones (2016) argues the weight of evidence suggests that people typically behave as though $\gamma > 1$.

4.4.1.4.2 Common Sense Check It may also be worth pausing, briefly, to reflect on whether the above findings are intuitively plausible. Common sense intuition does actually have some predictive power, when it comes to theories of human behavior, so we should be at least slightly wary of psychological claims that don't pass the smell test (Hoogeveen 2020).

First, can imagine living on the equivalent of $1,000 a year and being offered an opportunity to risk death for a ten-fold increase in your income. Now, you can imagine living on the equivalent of $100,000 a year and being offered the same opportunity. Our model and estimate for γ imply that a normal person would be less keen to take the gamble in the second case. At least speaking personally, this result strikes me as fairly intuitive.

As another informal check, we can reflect on the apparent historical implications of our parameter estimate. If $\gamma < 1$, then this would predict that impoverished people are willing to sacrifice *larger* fractions of their potential income to reduce risks to their lives. This prediction would be difficult to square with the conventional observation that demands for occupational safety seem to have risen dramatically over time (Costa 2004). Colonial enterprises in the early modern period, as an extreme case, apparently killed double-digit percentages of the European workers who chose to participate in them (Acemoglu 2001). If we assume that $\gamma < 1$, then it is more difficult to explain why jobs with this level of mortality risk would—one imagines—struggle to attract similar numbers of workers in present-day

Europe.

4.4.1.5 Implications of Sharply Diminishing Returns

Jones (2016) draws broad implications from the tentative finding that $\gamma > 1$. He suggests that we may gradually be transitioning into a world where increasingly large fractions of possible consumption are foregone for the sake of security and longevity. The steady rise in healthcare spending as a fraction of income, he argues, is one potential symptom of a more general trend. Increasing safety-oriented regulations and restrictions on innovation may be another. Aschenbrenner (2019) suggests that growing living standards may also exert upward pressure on spending in areas such as biosecurity.

Naively, a decline in the frequency or deadliness of war might be taken as another direct implication of this more general theory. Any material benefits that war can offer will become increasingly unimportant compared to the various leveling risks that it poses. Even if the potential gains from war grow in proportion with consumption, due to broadly distributed economic growth, this will seemingly not be enough to prevent the attractiveness of war from falling.

However, without additional analysis, we cannot yet arrive at this optimistic conclusion. As a next step, we should try to account for the significance of motives unrelated to consumption.

4.4.2 Accounting for Additional Incentives: A More General Model

4.4.2.1 Residual Goods

Consumption is a very general concept. A very wide range of things can be purchased, at least indirectly, and at least to some limited extent. Money can and often is spent on goods or services that help people form and maintain friendships, manage mental illness, enhance their own privacy, have spiritual experiences, and learn about the world. Models that represent utility as a function of consumption are therefore less simplistic than they

may initially appear. These models do not assume, for instance, that people only care about material comforts.

However, no one would deny that consumption levels are only a highly imperfect predictor of individual welfare. All purely consumption-based models of welfare will inevitably be associated with large residuals. *Residual goods* are factors that help to explain these residuals.

A number of residual goods frequently enter into individual decisions about whether to support or participate in wars: honor, moral satisfaction, social standing, and so on. Before we draw conclusions from consumption-based models of individual risk-taking, we should first consider the significance of these additional motives for war. Could taking them into account reverse any of our initial theoretical findings?

4.4.2.2 Gambles With Residual Goods

The simplest way to incorporate residual goods into our rational actor model is to treat them as an independent source of utility. This means that we can represent overall utility as the sum of utility from consumption, $U_C(C)$, and utility from residual goods, $U_E(E)$:

$$U(C, N) = U_C(C) + U_E(E) \tag{4.4}$$

Now, let us consider a gamble that offers both ΔU_E units of expected value from residual goods and a probability p of changing one's consumption by a factor of R.[19] For this case, we can also drop the assumption that $R > 1$. If the gamble offers enough residual value, then the individual may actually be willing to accept a pay cut in addition to some risk of death.

[19]Some gambles produce residual value even if the individual ultimately dies, faces imprisonment, or otherwise experiences a leveling event. For instance, a soldier may find it meaningful to fight and die for a good cause. For this reason, we have not – in this model – assumed that earning residual value requires avoiding the leveling event.

Adding ΔU_E to the left-hand side of expression (3) gives us the following inequality:

$$\Delta U_E + p[U(RC) - U(C)] > (1-p)[U(C) - L] \tag{4.5}$$

Let us suppose, in line with the evidence discussed above, that $\gamma > 1$. If $R > 1$, then the left-hand side of the inequality still declines as it did before. The only difference is that it now declines to ΔU_E rather than declining to zero. Meanwhile, as before, the right-hand side of the inequality rises to a maximum value of $(1-p)(A-L)$. The gamble will remain rational if and only if $\Delta U_E > (1-p)(A-L)$.

For $R < 1$, the left-hand side of the inequality actually increases. Although the gamble now involves sacrificing some portion of one's consumption, this sacrifice becomes less bothersome as one grows wealthier. The left-hand side therefore approaches ΔU_E from below. The net value of the gamble still decreases if $pR^{1-\gamma} < 1$, because the right-hand side grows more quickly than the left-hand side. The net value of the gamble instead increases if $pR^{1-\gamma} > 1$. In either case, as consumption rises, the gamble becomes or remains rational if and only if $\Delta U_E > (1-p)(A-L)$.

The core intuition behind these results is that, if $\gamma > 1$, then increasingly wealthy people care less and less about fractional changes in their consumption levels. Because there is an upper-bound on utility from consumption, there is also a cap on what they can stand to lose from leveling risks. The ultimate question for wealthy people, then, is whether the residual goods the gamble offers are worth the maximum leveling risk it can pose. For example, a wealthy person may focus on the question of whether it is worth risking death for the sake of status or for the sake of some moral principle. Both financial rewards and financial penalties simply become less relevant to them as their wealth grows.

4.4.2.3 The Minimum Risk Premium

We have just looked at how the expected value of a given gamble evolves, if we raise the initial consumption level and hold the other parameters fixed. It is also informative—

	$0 < \gamma < 1$	$\gamma > 1$
$\Delta U_E < (1-p)(A-L)$	Declines toward $\frac{1}{p}^{\frac{1}{1-\gamma}}$	Grows to become infinite
$\Delta U_E = (1-p)(A-L)$	Equals $\frac{1}{p}^{\frac{1}{1-\gamma}}$	Equals $\frac{1}{p}^{\frac{1}{1-\gamma}}$
$\Delta U_E > (1-p)(A-L)$	Grows toward $\frac{1}{p}^{\frac{1}{1-\gamma}}$	Shrinks toward 0

Table 4.3: The effect of rising consumption on the minimum value of R needed to render the gamble rational, holding other parameters fixed.

and perhaps ultimately more informative—to consider the minimum value of R necessary to render a gamble rational.

This value, R_{min}, reflects the difficulty of motivating individuals to fight. If R_{min} grows as consumption grows, then combatants will demand an increasingly large portion of available wealth. Compensating them will eventually become impractical. If R_{min} shrinks, on the other hand, then they will demand an increasingly small portion of available wealth. Compensating them will instead become increasingly easy.

Table 1 describes the evolution of R_{min} when p and ΔU_E are held fixed. Given available empirical evidence, the final column is most relevant.[20]

If we assume that $\gamma > 1$, and assume that p and ΔU_E are sufficiently small, then R_{min} will eventually blow up to infinity. If p and ΔU_E are sufficiently large, however, then R_{min} will be less than one. As a result, it will shrink over time, as income losses become increasingly painless to the individual.

4.4.3 Does Prosperity Actually Matter?

At this point, it is worth pausing to reflect on the fact that the above analysis is in tension with what might be called the *conventional view* of the relationship between prosperity and conflict. As noted above, a number of prominent scholars have dismissed

[20]Note that, if $\Delta U_E = 0$ and $\gamma > 1$, then R_{min} necessarily blows up to infinity. The reason is that A is an upper bound on utility from consumption when $\gamma > 1$. If A were less than L, then this would violate the assumption that the leveling risk represents a loss of utility.

the idea that prosperity levels are in-themselves relevant to combatant incentives (Fearon, 2008; Ray, 2017). Their suggestion is that, as prosperity rises, the upsides and downsides of conflict should change in ways that tend to counterbalance each. Structural changes (such as changes in the seizability of wealth) may influence combatant incentives, but combatant incentives do not change if everyone simply becomes richer.

The purpose of this section is to briefly stress-test our models and explore potential ways to support the conventional view.

4.4.3.1 Possible Objections

First, we can consider the objection Fearon (2008) makes to the relevance of prosperity. He suggests that, to support the view that prosperity matters, one must assume that individuals' utility functions do not display "constant relative risk aversion." However, we have not needed to make this assumption: the isoelastic utility function, which we have chosen to use in our model, *does* display constant relative risk aversion.

Second, on the flipside, we can consider whether the apparent relevance of prosperity could be an *artifact* of our decision to model the value of consumption using the isoelastic utility function. If one looks back at Section 4.4.1.3.2, though, one can see that our core analysis actually did not depend on any special properties of this function. The primary purpose of the isoelastic utility function, in this analysis, was simply to give us a convenient language (γ) for comparing the steepness of a person's returns curve to the steepness of a logarithmic function.

Third, we can consider the possibility that our findings depend on our decision to model prison and death as "leveling risks." We could instead have assumed that the downsides associated with imprisonment or death are constant: perhaps people's keenness to avoid these outcomes does not depend on how good their lives are. However, this alternative modeling assumption would have led to the same basic results. The net value of a gamble would have decreased if $\gamma > 1$ (because wealthier people have less to gain and the same

amount to lose) and have increased if $\gamma < 1$ (because wealthier people have more to gain and the same amount to lose).

Fourth, we can consider whether any special parameter values in our model could render prosperity irrelevant. We saw above that the minimum risk premium R_{min} will only be constant if $\Delta U_E = (1-p)(A-L)$. However, at least at first glance, there is no obvious reason to suppose that a knife's edge condition of this kind would be met for every person participating in every conflict. The condition also cannot be met if $\gamma > 1$, $A > L$, and $\Delta U_E \leq 0$. In other words, assuming $\gamma > 1$, the condition cannot be met if the individual fails to derive positive residual value from fighting.

4.4.3.2 Fearon's Alternative Model

Finally, we can attempt to understand better how Fearon (2008) arrives at the conclusion that prosperity is irrelevant. He presents a model that is, essentially, equivalent to our model except with three specific parameter value choices. First, the model assumes that $\gamma < 1$. Second, the model (implicitly) assumes that conflict offers no residual value (i.e. $\Delta U_E = 0$). Third, and most importantly, the model assumes that both death and imprisonment are equivalent in value to being alive with zero consumption (i.e. $U_C(0) = L = 0$, given our modeling convention that the value of being dead equals zero). This set of assumptions allows the knife's edge condition just described to be met, with both sides of the equation equaling zero.[21].

However, this set of assumptions does not strike me as reasonable. First, we have seen that it is probably false that $\gamma < 1$. If we instead assume $\gamma > 1$, then $U_C(0)$ blows up to negative infinity and, given the other assumptions, utility actually stops being mathematically well-defined. Second, it does seem reasonable to assume that conflicts do offer positive or negative residual value. Introducing this residual value would break the knife's edge. Third, even if keep the assumption that $\gamma < 1$—so that $U_C(0)$ remains mathematically

[21]Note that, when $\gamma < 1$, the assumption that $U_C(0) = 0$ implies $A = 0$

well-defined—I see no reason to assume that people would place equal value on being dead, being in prison, and being alive with zero consumption. The value of a life with zero consumption is not really conceptually well-defined, in any case, given that in practice people stop being alive once their consumption drops below the subsistence level. The function $U_C(C)$ should therefore only be thought of as referring to conceptually meaningful values for inputs greater than C_{sub}.

Although they do allow the knife's edge condition to be met, the assumptions made in this model are not attractive.

4.4.3.3 Summary

Ultimately, the claim that prosperity (in-itself) matters is robust. Only very weak modeling assumptions are needed to support it. Alternative models that imply that prosperity is irrelevant rely on implausible modeling assumptions. As noted above, the claim also seems to match common sense intuition and observation.

Although the *predicted strength* and even the *direction* of the prosperity effect depends on certain parameter estimates and modeling choices, which are comparatively open to debate, the basic conclusion that rising prosperity has some effect seems difficult to avoid.

4.4.4 Implications

This section has presented a simple argument regarding individual preferences. Since there are sharply diminishing returns on consumption, rising prosperity decreases the value people assign to the financial rewards that war offers. As a result, it becomes harder and harder for these rewards to compensate for the risk of death and imprisonment involved in civil war. All else equal, prosperity decreases incentives to fight.

Taking motives such as honor and moral satisfaction into account does somewhat complicate the picture. People who are sufficiently zealous and well-protected may favor conflict even if this means accepting a pay cut. When consumption levels rise, these people

actually become willing to forego increasingly large fractions of their potential incomes. In short, given our assumptions, risk-takers who *do not* require positive financial incentives to fight should become even easier to motivate. On the other hand, risk-takers who *do* require financial incentives should become more difficult to motivate.

Unfortunately, we have only very limited data on combatant income premiums. The conventional view is that wealth is often an important motivator for conflict participation, even if ideological motivations and other residual goods are important too (Collier 2004; Regan 2005; Blattman 2010; Ray 2017).[22] A plausible but tentative hypothesis is that, in a typical civil war, combatants *tend* to require some extra compensation beyond what they would otherwise be earning.

Therefore, tentatively, we should expect prosperity to lower the risk of civil war in most cases. By making it harder for rebel groups to attract fighters, in most cases, prosperity reduces the amount of fighting we should expect.

Ultimately, although the picture is much more complicated than it is initially seemed, the "Diminishing Returns Mechanism" does seem to be viable explanation for why states experience fewer civil wars as they become wealthier.

4.5 Other Relevant Empirical Evidence

If the Diminishing Returns Mechanism does in fact play a role in reducing the risk of civil war, as countries become wealthier, then there are a number of additional findings we might expect to observe.

[22]One notable study with a contrary finding is Bahney et al. (2013), which studies the compensation paid to Iraqi insurgents. The authors find that the salaries offered by al Qa'ida in Iraq were so low that it is unlikely that most insurgents with outside job options would have been earning meaningful wage premiums.

4.5.1 Risk Avoidance in Other Domains

First, we might expect to a similar pattern of rising risk-avoidance in other domains. People in wealthier countries should also be willing to forego larger amounts of income to avoid risks of death, imprisonment, or debilitating harm that come from sources other than war. Some relevant domains to consider are:

- **Healthcare spending:** We observe a similar pattern. Hall and Jones (2007) — who were a major source of inspiration for this paper — note that people in wealthier countries generally spend much larger fractions of their income on healthcare (i.e. on protecting themselves against medical risks).
- **Workplace safety:** We observe a similar pattern. Workplace safety expectations in developed countries are obviously both much stricter than they used to be and much stricter than workplace safety expectations in developing countries (LaDou et al., 2018). This is what we would expect if wealthier people are, in practice, less easily enticed to accept workplace risks.
- **Homicide:** We observe a similar pattern. Wealthier countries have, on average, substantially lower homicide rates (Ouimet, 2012).[23]

Of course, there may be alternative explanations for these observations. The homicide case, especially, has many possible explanations, since wealthier states tend (for example) to have more reliable policing and better social safety nets. Overall, though, these findings do at least somewhat bolster the hypothesis that the Diminishing Returns Mechanism is at work in the case of civil war.[24]

[23]Of course, some countries do have substantially higher or lower homicide rates than their wealth alone would predict. Famously, many countries in the Americas have anomalously high rates of crime (Bergman, 2018).

[24]These are the only three cases I have looked into. Since the first two came to my attention through the economics literature that inspired this paper, it would not be fair to consider them "randomly chosen." On the other hand, the third case was suggested to me as a possible counterexample. Overall, although this is not an entirely biased sample of cases, it is possible that a more random or diverse sample would turn up more cases that diverge from the pattern.

4.5.2 Insurgent Responsiveness to Shifts in Economic Incentives

Second, if the Diminishing Returns Mechanism is at work, then we should also expect to see evidence that insurgents are responsive to shifts in economic incentives. (If they were not, then shifts in incentives produced by rising prosperity would be irrelevant.) A number of studies, conducted over the past decade, support the notion that participants in insurgencies are in fact responsive to such shifts. These studies often examine the implications of changes in the values of particular commodities and particular kinds of labor (Dube, 2013; Berman, 2017; McGuirk, 2020; Blair, 2021).

For instance, Dube et al. (2013) find that increases in coffee prices are associated with declines in conflict in coffee-producing municipalities in Colombia. Meanwhile, increases in oil prices are associated with increases in conflict in oil-producing municipalities. They explain that, because the coffee industry is particularly labor-intensive, many potential combatants have the option of working in the coffee industry. Increases in coffee prices lead to increases in wages in the coffee industry, which lead to increases in the opportunity cost of violence, which lead to less violence. On the other hand, because the oil industry is very capital-intensive, increases in oil prices do not substantially increase the opportunity cost of violence. The primary effect is instead to increase the value of lootable assets. Hence, the argument goes, rising oil prices lead to more conflict. Insurgents at least appear to be responding somewhat rationally to shifts in economic incentives.[25]

[25]On the other hand, there are also studies whose findings are in tension with the hypothesis that insurgents respond to shifts in economic incentives. For instance, Berman et al. (2011) find that insurgent violence within Afghanistan, Iraq, and the Philippines is *negatively* associated with unemployment rates. We might expect to see violence increase as unemployment rates increase, since higher unemployment rates imply a lower opportunity cost to participating in insurgencies; men are less likely to need to forego economic opportunities in order to dedicate time to the insurgency. However, as the authors of the study note, this negative relationship may be explicable even if people are in fact responsive to incentives. For instance, violence may increase because insurgents need to use more violence to prevent economically deprived people from becoming paid government informants. The correlation could also reflect a causal relationship that runs in the opposite direction.

4.5.3 Insurgent Income Premiums

Third, if the Diminishing Returns Mechanism is at work, we should also expect to find that — at least in *most* conflicts — insurgents on average earn more income than they otherwise would. Our rational model suggests prosperity should only have a pacifying effect if people require compensation to be willing to take risks. (Otherwise, the effect may actually go in the opposite direction.) Unfortunately, there does not seem to be much data on insurgent incomes premiums.[26] Nonetheless, the limited information we do have does not clearly bear out this hypothesis.

As briefly noted above, Bahney et al. (2013) is one notable paper that finds that al Qa'ida in Iraq offered extremely low salaries, compared to the salaries reported in many postings for traditional low-skill jobs in Iraq. This observation suggests that members of al Qa'ida probably were not earning meaningful wage premiums. If typical members of al Qa'ida were in fact taking paycuts, then that should give us pause. On the other hand, it is important to remember that our model does allow insurgents to rationally accept paycuts when they are sufficiently strong ideologically motivated; for growth to tend to produce a decline in conflict, it only needs to be the case that insurgents in *most conflicts* are unwilling to take pay cuts. It remains possible that, even if the typical person took a paycut to join al Qa'ida, that this case was somewhat unusual due to a higher-than-average level of ideological motivation among members of al Qai'da.

4.5.4 Summary

Ultimately, this brief exploration of secondary evidence does not seem like it should dramatically shift the credence we give to the Diminishing Returns Mechanism. We have noted some supportive results, but the additional support they provide is arguably not very

[26]The scarcity of data is likely due to the difficulty of collecting it: relevant difficulties include accounting for the fact that compensation may come in large part through non-salary benefits (i.e. housing) or access to captured wealth, judging what someone's income *would be* if they were not a member of an insurgent group, and, more generally, gaining enough access to insurgents and insurgent groups to learn the relevant information.

strong. We have also noted some results that raise questions, but these results also do not seem to constitute strong debunking evidence either.

4.6 Conclusion

4.6.1 Recapitulation

Countries tend to become more peaceful as they grow wealthier. Scholars of civil war have put forward many different causal mechanisms to explain this relationship. I expect that several of these causal mechanisms do, in fact, play important roles in driving down the risk of conflict.

However, there is another (largely neglected) mechanism that may also play an important role. Specifically, economic growth may reduce the risk of civil war by making potential combatants more satisfied with their current income levels and, therefore, less willing to risk dying or being locked up for the sake extra income.

Simple rational actor models and empirical evidence concerning diminishing returns on consumption support the existence of this Diminishing Returns Mechanism. It is also supported by some secondary empirical findings, for example concerning the relationship between wealth and risk acceptance in other domains. On the other hand, there are some theoretical considerations and empirical findings that point in the other direction. For example, in the case of highly ideologically motivated combatants, our models suggest that growth may actually *increase* the risk of conflict.

Overall, though, the Diminishing Returns Mechanism provides a plausible explanation for the finding that growth (seemingly) reduces the risk of civil war.

4.6.2 What About Interstate War?

It seems worth asking whether the Diminishing Returns Mechanism might also operate in the case of *interstate* war.

Financial incentives presumably do influence the behavior of people involved in interstate wars too. Nothing about the mechanism is obviously specific to civil war; we have even discussed how it may apply in contexts as diverse as healthcare spending and occupational safety. Admittedly, the case of interstate war does seem more complicated, in part because the relationship between individual preferences and the chance of two states fighting seems comparatively indirect. However, the case does not seem fundamentally dissimilar.

If the Diminishing Returns Mechanism *does* operate in the case of interstate war, then that would give us at least one reason to suspect that growth lowers the risk of interstate war too. Interestingly, unlike in the civil war case, there is no academic consensus that growth makes interstate war less likely. One way to improve our understanding of the relationship between economic growth and interstate war, therefore, may be to try to understand whether and how the Diminishing Returns Mechanism applies in this more complicated case.

Chapter 5

Growth and Interstate War

5.1 Abstract

Economic growth probably lowers the risk of interstate war. This hypothesis is supported by various theoretical arguments. One of them is: Since there are strongly diminishing returns on wealth, growth tends to erode financial incentives to fight. Growth also has many downstream effects — for example, making states more economically interdependent and making it harder to seize states' primary forms of wealth — that probably contribute to peace as well. I clarify and synthesize these arguments within a formal bargaining model. Along the way, I also note a number of pathways by which growth can instead raise the risk of war. Finally, I argue that statistical evidence for a causal link between economic development and interstate peace is much stronger than is typically recognized. When we fit fixed effect models to conflict data, pairs of states seem to become about a third as likely to fight whenever their incomes double. Since widespread growth is a very modern phenomenon, it may at least in part help to explain an apparent post–WWII downward trend in the rate of conflict. However, there are also reasons to expect its effect to become smaller over time.

5.2 Introduction

Up until the nineteenth century, interstate wars generally took place against a backdrop of extreme scarcity. Preindustrial societies were almost universally poor (Landers, 2005; Clark, 2008). While some societies were more prosperous than others, or were more or less prosperous at different points in time, they rarely produced many times more output than was needed to keep their populations alive. $1,000 per person was probably close to a typical level of annual output.[1]

When we look back on pre-industrial wars — ranging from the Seven Years' War in the eighteenth century, to Sargon's conquests in the third millennium BC — we mainly see materially deprived people fighting to control very limited pools of wealth (Crone, 2015).

To offer one concrete illustration, the Hundred Years War — which England waged against France across the fourteenth and fifteenth centuries — was an unusually extractive war. English forces managed to capture the King of France and negotiated a financially ruinous ransom in exchange for his return. They also gained control of nearly than half of France's territory, for many years, and extracted the wealth they could from it. Nonetheless, in modern terms, the conflict probably only brought in around $1 billion in revenue across its entire decades-long run.[2] As a present-day point of reference, $1.5 billion is the approximate *annual* revenue of British McDonald's franchises (Global Database, 2021). France simply did not have much wealth to take.

At the same time, owing to England's own poverty, most of the English people who exposed themselves to terrible risks by fighting in the war did not need to be paid very

[1] Maddison (2007) estimates that the average global income in 1500AD was equivalent to $1,300 in 2020 US dollars. However, this specific number should be interpreted as only an educated guess rather than a tightly data-driven estimate.

[2] Postan (1964) estimates that war indemnities brought in around £400,000 and that the most notable ransoms brought in around £250,000 (not adjusted for inflation). He regards other sources of revenue as more difficult to estimate, but unlikely to lead to a far greater overall figure. Speculatively, then, we can use £1,000,000 as a rough estimate for the overall revenue brought in by the war. The Bank of England's inflation calculator suggests that this sum is roughly equivalent to £900,000,000 in present-day pounds. Converting this into dollars yields an estimate just slightly above $1,000,000,000. Of course, this particular number is probably not very accurate. Still, we might expect it to be at least within the correct order-of-magnitude.

much. The longbowmen who played an essential role in supporting English victories earned wages roughly equivalent to $4,000 per year, which was nonetheless far more than they would otherwise have been earning.[3] For reference, an individual earning this income in modern Britain — even after their housing costs are subtracted — would be classified as falling about 30% short of the line for no longer being considered "destitute" (Fitzpatrick et al., 2023).

States only began to meaningfully rise out of this kind of poverty in the nineteenth century, as a result of the Industrial Revolution (Stearns, 2020). Substantial, sustained economic growth emerged in Britain and then spread outward. By the second half of the twentieth century, most states had begun to grow. (See Figure 5.1.[4]) Today, all but the least developed states in the world are probably more prosperous than any pre-nineteenth-century state.

Many economic historians see the emergence of modern economic growth — the start of the so-called "Great Enrichment" — as a central pivot point in human history (e.g. Clark, 2008; Van Zanden, 2009; Galor, 2011; McCloskey, 2016). For instance, McCloskey (2004) writes: "It is certainly the most important event in the history of humanity since the domestication of animals and plants, perhaps the most important since the invention of language."

In contrast, histories of global politics seldom emphasize the emergence of growth.[5] Nonetheless, it is natural to suspect that the transition from a world of scarcity to a world of relative abundance has impacted the interactions of states. The absence of such an impact

[3]This estimate was generated by taking figures reported by Taliaferro (2015) and again using the Bank of England inflation calculator. Estimates reported in the same source also suggest that aristocrats who made up a minority of English forces were of only modest means: knights' wages were roughly equivalent to $30,0000.

[4]To generate this graph, I have filled in gaps in the estimates for individual countries by assuming smooth growth within the gaps. I have also extrapolated countries' initial estimated values back into earlier years; I expect this extrapolation approach to help us more accurately identify the median country in a given year, compared to the approach of simply excluding these (disproportionately poor) countries from consideration in a given year.

[5]Buzan and Lawson (2015), for instance, argue that the field of International Relations has largely ignored the importance of the economic transformation of the world over the course of the nineteenth century. However, even their book-length treatment of this transformation does not devote significant space to discussing rising prosperity. They instead focus on issues such as growing regional inequality and global interconnectedness. Gat (2017), who does highlight the Great Enrichment in his account of the history of war, is notable for being a rare exception.

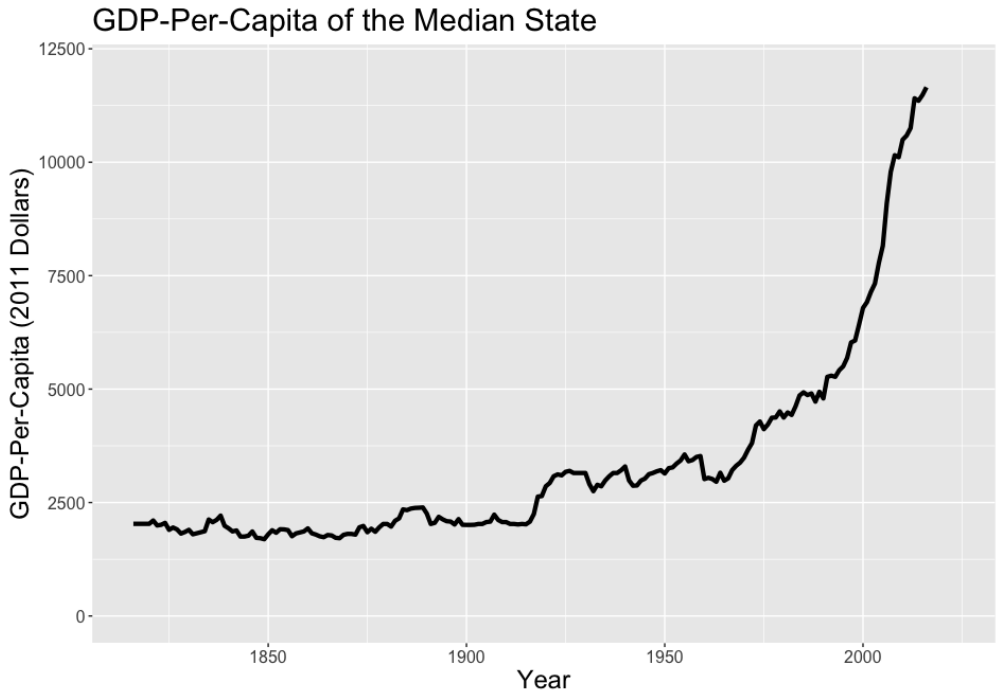

Figure 5.1: The purchasing-power-adjusted GDP-per-capita of the median state over time, for states included in the Correlates of War dataset. GDP-per-capita estimates are derived from the 2020 Maddison Project Database (Bolt and van Zanden, 2020). We can see that the median state began to experience substantial economic growth only in the mid-twentieth century, although Britain began its own growth take-off more than a century earlier.

would be interesting in its own right.

This paper argues that this growth has probably had an overall pacifying effect. Because there are strongly diminishing returns on consumption, it is harder for financial rewards to motivate affluent people to tolerate the risks wars pose. The richer someone is, the less inclined they will be to risk death or internment for the chance to multiply their income. Even if financial rewards grow in proportion to the individual's baseline income, due to rising national budgets or expanding pools of wealth to plunder, this will not be enough to offset the significance of strongly diminishing returns. Therefore, prosperity naturally tends to dampen individuals' financial motives to participate in or support wars. This shift probably, in turn, has a pacifying influence on the foreign policy decisions states make.

Of course, when states undergo economic growth, they do not simply become more prosperous. States that grow also tend to experience a wide range of structural economic

changes. Their sources of wealth tend to shift from land and other natural resources to a variety of more complex forms of capital. They become more intertwined with other states through trade and investment relationships. Their internal markets become more active and complex. And so on. Some of these downstream, structural implications of growth probably support peace as well — and may even be more important. Growth probably exerts pacifying effects through multiple channels.

The first section of the paper summarizes previous theoretical writing on the relationship between growth and interstate war. Many other scholars have discussed the significance of structural implications of growth, especially greater economic interdependence. However, surprisingly, there has been very little writing on the significance of prosperity itself or the implications of the diminishing returns people derive from consumption.

The second and third sections develop and comment on a formal model that relates prosperity to the risk of interstate war. The model shows that total wealth influences the set of peaceful arrangements that two states would find mutually beneficial. The larger this set becomes, the more likely the states are to avoid war. I show that—if we assume the returns on consumption diminish sufficiently quickly—then prosperity lowers the risk of financially motivated wars. As people become richer, they become less motivated to fight for the sake of extra money.

However, of course, people can also have non-financial motives to fight. This observation suggests a distinction between *scarcity-driven wars*, which are rational in part due to scarcity, and *scarcity-constrained wars*, which would be rational even if no one cared about their income. Prosperity actually exacerbates the risk of scarcity-constrained wars by making people more willing to suffer income losses in order to fight. Fortunately, this aggravating effect is probably comparatively small. At least for sufficiently poor states, prosperity probably still lowers the overall risk of war.

The fourth section explores other, downstream implications of economic growth beyond rising prosperity. I survey other pathways by which growth can support peace—for

instance, by fostering economic interdependence—and relate them to the bargaining model developed above. I also note a number of additional pathways by which growth can support conflict. The section is, in essence, an attempt at synthesizing the theoretical "Capitalist Peace" literature.

The fifth section presents empirical evidence that pairs of states become much less likely to fight each other when they grow wealthier. This latter observation has been underappreciated by international relations scholars, perhaps due to the field's historical wariness toward the use of fixed effect models. We have fairly strong empirical grounds to suspect that growth has been a force for peace. Rising interdependence also seems to account for at most a small portion of the relationship between growth and peace.

I conclude by summarizing the paper and offering several caveats. Unfortunately, even if growth has tended to support peace, this does not necessarily imply that the risk of war is in decline. Other factors could outweigh the pacifying effects of growth. Furthermore, since growth can have both aggravating and pacifying effects, we cannot even be confident that growth will continue to be pacifying on net.

5.3 Previous Work

There is a large and long-standing "Capitalist Peace" literature discussing how economic variables may influence the risk of interstate conflict (e.g. Angell, 1913; Gartzke, 2007; Schneider and Gleditsch, 2010; Schneider, 2017). I will save a more detailed overview of this literature for Sections 5 and 6 of this paper, but make a few key points now.

First, the literature has largely focused on structural or qualitative aspects of national economies and the global economic system. For example, the effect of trade and economic interdependence between states has been very widely discussed. The effect of having an economy based on industrial production, the effect of receiving foreign investments, and the effects of several other variables have also been discussed a fair amount.

Since economic growth tends to involve or support a wide range of qualitative economic changes, a substantial portion of this literature implicitly speaks to the effect of economic growth. (Section 5 of this paper will provide a more detailed overview of the relevant literature and its implications of the effect of growth.) However, growth has not typically been a focus of the literature. As Gartzke (2015) has written, growth (or "development") is the "neglected stepchild of the study of economics and conflict."

Beyond commenting on qualitative implications of growth, the literature has said surprisingly little about the direct significance of countries becoming less poor. Gat (2017) is, to my knowledge, the most significant work to note the potential relevance of prosperity itself on the risk of interstate war. Gat argues that, because there are diminishing returns on consumption, prosperity should dampen financial incentives to fight. However, Gat does not formalize his argument or clarify the assumptions it hinges on.[6]

Outside of the interstate war literature, in a discussion of the causes of *civil war*, Fearon (2008) also raises the idea that, due to diminishing returns, prosperity reduces the financial incentive to become an insurgent. However, he then rejects the idea on the basis of a formal model in which prosperity has no effect. In the paper that made up the previous chapter of this thesis, I argued that this result hinged on implausible modeling assumptions and that it is hard to escape the conclusion that prosperity matters. (The core model developed in that paper serves as a foundation for the more complex interstate war model this paper will develop.)[7]

One reason that the interstate war literature has paid relatively little attention to economic growth is that — despite theoretical arguments suggesting that growth *should* have an effect — a number of classic studies have failed to report a robust statistical association between the level of economic development of a pair of states and the chance that they will

[6]For example, Gat (2017) assumes that people experience "diminishing returns." However, as we will see below, this assumption is not actually strong enough to support the argument. Returns need to diminish *quickly enough* to compensate for the fact that combatants in wealthier societies can also hope to earn more money.

[7]As in that paper, this paper takes significant inspiration from theoretical research on the behavioral implications of growth (Hall and Jones, 2007; Jones, 2016; Aschenbrenner, 2019).

fight (e.g. Maoz and Russett, 1992; Mousseau et al., 2003; Gartzke, 2007). For instance, in one review article on the economics of conflict, Schneider (2017) states, "[D]yadic development does not reduce the probability of interstate war." Scholars arguing for the existence of an association have largely focused on interaction effects, for instance between development and democracy (Mousseau, 2000; Mousseau et al., 2003) or between development and contiguity (Gartzke, 2007). Some scholars have also suggested there is evidence for "spillover effects," whereby economic development in some states lowers the probability that other less-developed pairs of states will fight (Gartzke and Weisiger, 2014). This paper will ultimately diverge from most of this empirical literature by suggesting that there actually *is* a strong association between development and peace between states, at least once controls for fixed traits of pairs of states are added to the relevant models.

5.4 Growth and Interstate War: The Diminishing Returns Mechanism

5.4.1 A Bargaining Model

5.4.1.1 Core model

In this section, I develop and explore a formal model that relates wealth to the risk of interstate war. The chief point of the model is to clarify how broadly distributed growth can influence the risk of conflict, even in the absence of any structural economic changes. Since "mechanism of action" in the model depends on how strongly people experience diminishing returns on consumption, we can say that the model describes the *Diminishing Returns Mechanism* linking growth to the risk of conflict.

Like all formal models, the model considers a highly idealized case. Nonetheless, I expect that many of its idealizations can be relaxed without altering its qualitative predictions. (I will talk more about the significance of some of its idealizations below.)

At a high level, the model describes two states that are acting to maximize the welfare of their people. A war between them would kill people and destroy wealth, but would redistribute part of the remaining wealth from the loser's surviving population to the victor's surviving population. The people in these states may also have non-financial reasons, for instance ideological reasons, to be attracted to or averse to fighting. The model looks at how wealth levels influence the likelihood of these two states avoiding war.

The model's features are as follows:

- There are two states that might go to war with each other. If they do, then their respective probabilities of "winning" the war are p_1 and $p_2 = (1 - p_1)$.

- Before the war, wealth-per-capita across the two states is W. After the war, the winner will possess some fraction α of this wealth and the loser will possess some fraction $\beta < \alpha$. Since some portion of the wealth μ will be destroyed over the course of the war, it is the case that: $\alpha + \beta = (1 - \mu) < 1$.

- In expectation, some fraction v of the victor's population will survive. In expectation, some fraction d of the defeated state's population will survive. For convenience, we can also define the variables $m_1 \equiv 1 - vp_1 - dp_2$ and $m_2 \equiv 1 - vp_2 - dp_1$. These describe the expected mortality rate for each state, averaging across both possible outcomes.[8]

- The utility each person derives from wealth is given by the *isoelastic utility function*: $U(W) = A + \frac{W^{1-\gamma}}{1-\gamma}$. This function, commonly used in economic models, makes it possible to represent a range of returns curves by adjusting the value of γ. When $\gamma = 0$, the utility function is linear and the individual experiences no diminishing returns on consumption. As $\gamma \to 1$, the utility function increasingly approximates a logarithmic utility function. For $\gamma > 1$, returns on consumption are sublogarithmic. The parameter A is simply a constant term. Without loss of generality, we can associate the state of being dead with 0 units of utility.

[8]For simplicity, this model focuses on the risk of death and ignores other non-financial risks such as imprisonment or injury.

- If the war occurs, then each person in each state i also receives a utility "bonus" or "penalty" of B_i. This represents the utility gained or lost due to non-financial factors such as pride, excitement, moral satisfaction, social status, trauma, and shame. People receive the bonus or penalty even if they ultimately die.

- Within each state, everyone has the same level of wealth and exposure to risk. The populations of the two states are also equal.

- The leaders of both states are perfectly responsive to the preferences of the people they contain. Specifically, the leaders of both states act as though they are attempting to maximize average welfare.

5.4.1.2 Bargaining ranges and the risk of war

In this model, war will typically be *economically inefficient*. In other words, because war involves death and the destruction of wealth, there will typically be some peaceful division of the initial wealth that would leave both communities better off than if they went to war. The range of possible divisions that satisfy this property constitutes the *bargaining range*.

Rational pairs of states should only go to war, then, if they fail to identify and enact some division of wealth that falls within the bargaining range. In many cases, avoiding war may be straightforward, since both parties will believe that the "status quo" division is already preferable to war. For example, presumably, the United States and Canada would be happier to simply maintain their present division of wealth than they would be to go to war.

In other cases, states may need to actually make some substantial concessions to avoid conflict. In their bluntest form, concessions can take the form of explicit payments or tributes. Although these are less common today, they were quite common pre-modern international relations. For instance, in the context of the aforementioned Hundred Years' War, there

were multiple points where English and French forces engaged in explicit bargaining about wealth transfers (Tuchman, 1987). Ahead of the Battle of Poitiers, as one example, there was flurried discussion of avoiding the battle in exchange for an English commitment to transfer money, land, and prisoners to France and refrain from attacking it again for at least seven years.

Concessions can also take much subtler forms. For example, a weaker state can adopt policies that are favorable to the dominant state, motivated in part by an unspoken threat of future violence. This kind of "unspoken bargaining" to avoid conflict is probably, in practice, more prevalent than the explicit kind.

As Fearon (1995) has shown, however, the existence of a non-zero bargaining range does not guarantee that even fully rational pairs of states will avoid war.

First, an absence of shared information could lead them to disagree about the values of key parameters. For instance, because they lack information about each other's capabilities, they may both overestimate their own probabilities of victory. As a result, they may then have entirely inconsistent beliefs about where the bargaining range actually lies — and therefore be unable to identify any division of wealth that they both believe would actually be preferable to going to war.

Second, even if the actors agree that some bargain would be mutually beneficial, they may not be able to trust one another to abide by this bargain. For instance, a state whose power is declining may worry that a rising power will attempt to revise the division of wealth once its position improves. In the aforementioned Poitiers case, as another concrete illustration, it would have been very reasonable for the French to wonder whether English were actually going to honor their side of the proposed bargain: perhaps they would not actually hand over everything they promised, or perhaps they would actually attack again (with greater strength) within the seven-year window and try to take some of their concessions back.

Fortunately, the larger the bargaining range is, the weaker these barriers to peace

become. Actors are more likely to agree that at least one bargain is mutually beneficial if a larger fraction of bargains actually are mutually beneficial. Larger bargaining ranges will also tend to contain some bargains whose expected values, conditional on compliance, are especially high relative to the expected value of going to war. The larger this difference in expected value is, the more willing rational actors will be to tolerate some risk of non-compliance.

Furthermore, when some possible war is already irrational, increasing the size of the bargaining range will tend to decrease the *minimum degree of rationality* that is sufficient for the two communities to avoid war. The decision to fight will become more and more obviously foolish. The basic line of argument here still applies, therefore, even if we do not assume that states are in fact very rational.[9]

5.4.1.3 How prosperity changes bargaining ranges

For the model just described, we can describe the long-run effect of rising wealth on the size of the bargaining range. Table 5.1 summarizes the effect for different sets of parameter values. The values are normalized so that a bargaining range of size zero would imply that no possible division of the initial wealth is preferable to war, while a bargaining range of size one would imply that all possible divisions of the initial wealth are preferable to war.

I will focus on the rightmost column, since empirical evidence tends to suggest that $\gamma > 1$ (Jones, 2016).[10] This means that people tend to experience sub-logarithmic returns

[9]Although the specific bargaining model we have presented directly represents financial and immaterial (e.g. ideological or honor-based) motives for conflict, it does not directly represent *security motives*. Of course, however, security motives can play an important role in informing decisions to go to war: states may sometimes go to war at least partly because they are concerned about other states attacking them later on. (Thucydides famously listed "fear" as one of the three central motives for war, along with "honor" and "self-interest.") However, it seems reasonable to think of security motives as largely downstream of other motives. If two states both have strong incentives to avoid conflict with each other, then, beyond a certain point, mutual fear seems hard to sustain. England and France are not currently worried about attacks from each other, for example, in part because these attacks would currently make no underlying sense. For this reason, I generally expect security motives to also grow weaker as the bargaining range expands.

[10]See the previous chapter of this thesis for a much more extensive discussion of the relevant evidence.

	$0 < \gamma < 1$	$\gamma > 1$
$(B_1 < m_1 A)$ and $(B_2 < m_2 A)$	Shrinks to $1 - 2(p_1 v^\gamma \alpha^{1-\gamma} + p_2 d^\gamma \beta^{1-\gamma})^{\frac{1}{1-\gamma}}$	Expands to 1
$(B_1 > m_1 A)$ or $(B_2 > m_2 A)$	Expands to $1 - 2(p_1 v^\gamma \alpha^{1-\gamma} + p_2 d^\gamma \beta^{1-\gamma})^{\frac{1}{1-\gamma}}$	Shrinks to 0

Table 5.1: The effect of rising wealth on the size of the bargaining range, for different sets of parameter values. A bargaining range of 1 implies that any possible division of wealth is mutually beneficial relative to war. A bargaining range of 0 implies that no possible division of wealth is mutually beneficial relative to war.

on consumption. When $\gamma > 1$, there is a ceiling on how much utility people can derive from consumption and A represents the value of this ceiling.

If $B_i < m_i A$ for both states, then, this means that sufficiently well-off people do not consider the non-financial *bonus utility* the war offers large enough to outweigh the risk of death. In other words, if the war had no financial implications, they would be opposed to it.

If $B_i > m_i A$, on the other hand, then the bonus utility alone is sufficient to outweigh the risk of death. In other words, if the war had no financial implications, people would support it.

The table reports that wealth expands the bargaining range (decreasing the risk of war) in the first case. Meanwhile, wealth shrinks the bargaining range (increasing the risk of war) in the second case. What explains these results?

The basic driver of the results is that, if people derive sub-logarithmic returns from consumption, then they become less sensitive to proportional changes in their incomes as they become wealthier. For example, if someone derives sub-logarithmic returns on consumption, then this means that they care more about going from a $1,000 income to a $10,000 income than they care about going from a $10,000 income to a $100,000 income.[11]

It then follows that, as the overall amount of available wealth increases, people become less sensitive to what *portion* of this wealth they hold. For example, the difference between holding 40% or 50% of a pool of wealth matters less when the pool is very large. The states in

[11] In contrast, if they derived *super-logarithmic* returns from consumption, then they would instead value the second income increase more highly. For example, if they had a linear utility function—meaning that they did not experience diminishing returns—then they would assign ten times as much value to the second increase.

our model therefore care less about the international division of wealth as the overall amount of wealth in the system increases. Even if war could lead to a more favorable division of wealth, welfare-maximizing states will become less and less willing to kill off portions of their populations for the sake of wealth redistribution. Therefore, if the states' populations only have financial motives to support a war, then rising wealth will tend to lower the war's probability.

The picture becomes more complicated, though, once we introduce non-material motives to fight. If the non-material value the war would offer these populations (B_i) is significant enough to outweigh the risk of death they would face (m_i), then the effect of wealth can actually reverse. When $B_i > m_i A$, financial disincentives are the only thing that can prevent war from being rational. People are inclined to fight, for non-material reasons, but may ultimately be discouraged by a fear that wealth destruction or unfavorable wealth distribution will lead them to become poorer. The wealthier they are, though, the less they will actually care about losing some portion of their wealth. Financial disincentives to fight will become less and less important, until eventually there is nothing stopping war from breaking out.

This explains why, according to the model, rising wealth lowers the risk of war when $B_i < m_i A$ but raises the risk of war when $B_i > m_i A$.

5.4.1.4 Exploring effect sizes

Of course, even if growth has some impact on the risk of war, we cannot necessarily assume that the effect will be large. In the context of the model just analyzed, the effect size is highly sensitive to various parameter values. Various plausible extensions of the model – for instance, extensions that acknowledge that not all individuals within a community are placed equally at risk, that acknowledge the rather more complicated relationship between individual preferences and the preferences of key decision-makers, that acknowledge the psychology of loss aversion, and so on – should also be expected to predict fairly different effect sizes, even if they continue to agree on the direction of the effect. Furthermore, it is

not clear how to convert shifts in bargaining range sizes into shifts in the probability of war. All of this leaves open the possibility of an effect that is real but trivial in its magnitude.

Nonetheless, this particular model does suggest that large effect sizes are plausible at least for sufficiently destructive wars. As an illustrative case, let us consider the following hypothetical war: The victor can be expected to lose .1% of its population and control 50% of the initial wealth, the defeated community can be expected to lose 1% of its population and control 40% of the initial wealth, and both communities are equally likely to be victorious. People in both countries would derive some non-consumption value from the war: no matter how high their consumption level rises, the typical person would be willing to accept at least a .1% chance of death in exchange for this additional value. People are fairly risk-averse: $\gamma = 1.3$. Finally, we need to choose a value for A: We can pin it down by supposing, speculatively, that a typical person with a million dollar income will accept at most a 5% probability of death for the chance to increase their income to an indefinite height.

In this case, the bargaining range will increase by approximate 35% as incomes rise from the approximate average level in 1500AD ($1,000) to the approximate average global level today ($20,000). See Figure 5.2 below for a graphical depiction of the effect of rising wealth. If another 200 years pass, with an average annual growth rate of 3%, then the bargaining range will ultimately become nearly four times larger than it would have been in 1500AD. These expansions of the bargaining range, especially the predicted long-run expansion, are fairly non-neglible.

Tweaks to the parameters can, however, lead to quite different predictions. If we decrease the expected mortality rate for each country by factor of ten, then growth produces a (very slowly) shrinking bargaining range. See Figure 5.3. On the other hand, if we then re-adjust the mortality rates to a level more consistent with nuclear war, with 25% and 50% of the people in each country dying, then growth has a dramatically pacifying effect. See Figure 5.4.

Our model suggests, ultimately, that growth has the best chance of reducing the

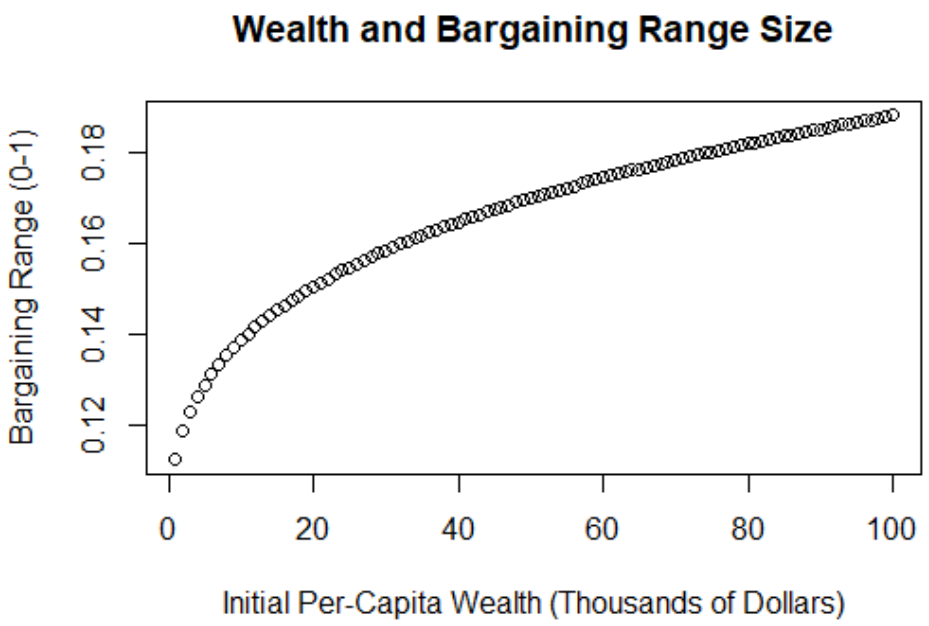

Figure 5.2: The effect of rising wealth on the size of the bargaining range. The relevant parameters are B_1= .001, B_2=.001, γ =1.3, v = .999, d = .99, α = .5, β = .4, p_1 = .5, and A = 1.07.

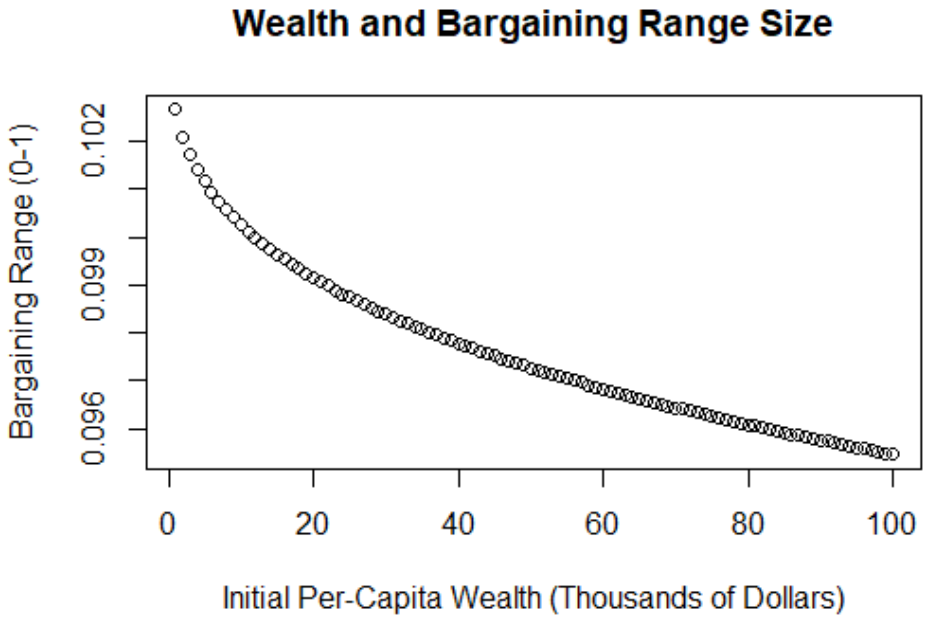

Figure 5.3: Reducing the fatality level can flip the sign of the effect. Note, though, the small magnitude. The relevant parameters are B_1= .001, B_2=.001, g =1.3, v = .9999, d = .999, α = .5, β = .4, p_1 = .5, and A = 1.07.

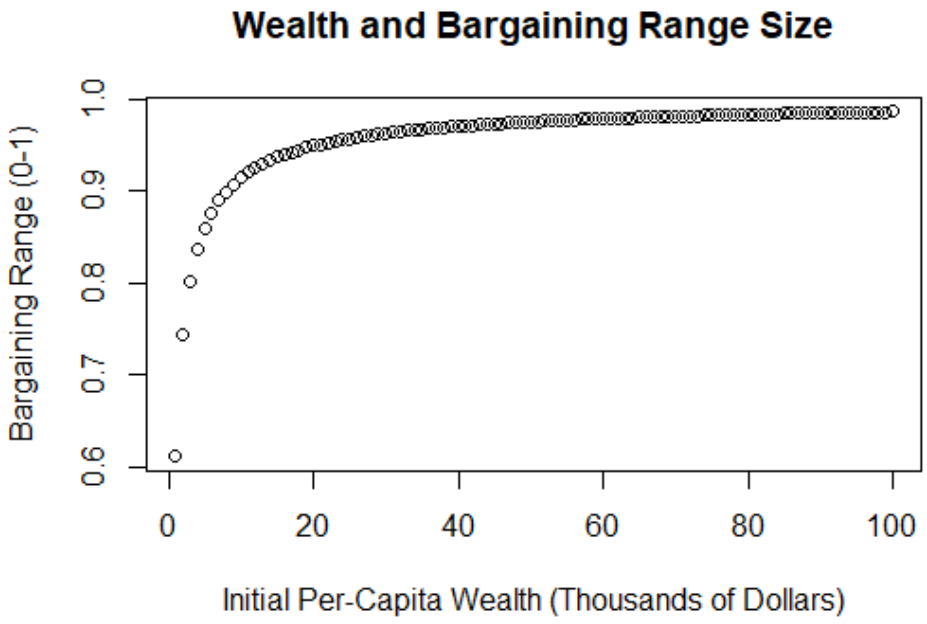

Figure 5.4: For very high fatality levels, the effect of rising wealth can be dramatic. The relevant parameters are B_1= .001, B_2=.001, g =1.3, v = .75, d = .5, α = .5, β = .4, p_1 = .5, and A = 1.07.

frequency of conflicts that are foreseeably very deadly.

5.4.1.5 Summary

We have explored a bargaining model that relates the prosperity to the risk of interstate conflict. We found that, depending on the details of a conflict, prosperity can have either a pacifying or aggravating effect. Overall, prosperity is more likely to have a pacifying effect when a conflict is associated with a high expected mortality rate and offers only limited non-material benefits.

However, our model made a number of notable simplifying assumptions. It also left open the question of whether we should expect prosperity to raise or lower the risk of war *overall*. The next section will relax some of these assumptions and attempt to clarify prosperity's overall effect.

5.4.2 Further Considerations

In this section, I will discuss a number of important issues that are either ignored or raised by the bargaining model. Ultimately, the primary purpose of this section is to bolster the hypothesis that rising prosperity has had an overall pacifying effect.

5.4.2.1 How do individual preferences influence the risk of war?

Our bargaining model assumes that states act as though they are trying to maximize the welfare of their populations. This assumption allows us to draw a direct link between changes in the incentives of at-risk individuals and the incentives of states.

Of course, though, the assumption is not exactly true. The relationship between individuals' incentives and the incentives of states is rather more complex. Nonetheless, we can still describe a number of pathways by which rising prosperity can influence states' inclinations to fight. For the sake of simplicity, I will ignore non-material incentives for conflict and focus here on pacifying effects.

First, as prosperity increases, it may become increasingly costly to entice people to accept high-risk positions in military organizations. Leaders might then become more averse to initiating violent conflicts, because they recognize that, if they initiate these conflicts, they will face more difficult problems with recruitment, retention, and combat effectiveness. Particularly in authoritarian countries, they may have greater reason to fear military coups or uprisings supported by disaffected former members of the military.[12]

Second, even if a war produces economic benefits, it may become increasingly difficult to win broader public support for the war by redistributing these benefits. Leaders may then have greater reason to fear disruptive popular uprisings, loss of power (by either democratic or non-democratic means), or lower levels of effectiveness (due to the various ways in which civilians can support a conflict).

[12]Notably, even states that use conscription also tend to use financial inducements to some degree. This suggests that states believe that financial inducements are an important tool for achieving good outcomes once a war has been initiated.

Third, in some cases, it is actually reasonable for leaders' and other elites to fear death or imprisonment. Sufficiently intense interstate conflicts may provoke assassination attempts, result in the capture and execution of opposing political elites, or escalate to the level of nuclear war. If a war would have produced financial gains for a state's elites, which can happen in sufficiently corrupt states, then elites' incentives to support a war may also decline as they become wealthier.

Fourth, it may become more difficult for war-prone leaders to rise to power or retain it. For instance, other things being equal, voters in democratic states may be increasingly unwilling to vote for leaders who are perceived to have an unusually high chance of provoking high-fatality interstate wars. The same consideration holds for any leader whose power is dependent on the support of groups that would be placed at risk in war.

Finally, leaders may simply care about the welfare of their citizens to some degree. They may also intrinsically value their national popularity or popularity with particular groups. More speculatively, leaders' orientations toward war may be influenced by broadly shared cultural norms within their societies, which might in turn be influenced by rising prosperity. As people become less willing to risk their own lives, and observe this same unwillingness in others, perhaps it becomes increasingly socially unacceptable to place lives at risk. Both leaders and members of the broader public may then internalize these norms and act in accordance with them.

In summary, even if we drop the model's simplifying assumption that states are purely focused on maximizing welfare, we should continue to expect prosperity to influence the decisions states make. This section has outlined a number of pathways through which rising prosperity may lower the risk of war. Of course, if we take non-financial motives for war into account, then it is also easy to sketch equivalent pathways through which war could have an aggravating effect.

5.4.2.2 Scarcity-driven wars and scarcity-constrained wars

5.4.2.2.1 Two types of war Building on our bargaining model, we can draw a distinction between two different kinds of wars.

Scarcity-driven wars are wars that are made more likely by scarcity. If everyone was wealthier, then these wars would be less likely to happen. In our model, of course, wars are scarcity-driven when $B_i < m_i A$.

Scarcity-constrained wars are wars that are made less likely by scarcity. If everyone was wealthier, then these wars would be *more* likely to happen. In our model, wars are scarcity-constrained when $B_i > m_i A$.

In this section, I discuss the relative prevalence of these two kinds of wars and how it may change over time. I tentatively conclude that prosperity has probably had an overall *overall* pacifying effect, over the past few centuries, even though it only makes one variety of interstate war less likely.

5.4.2.2.2 The Russo-Ukrainian War Of course, for any given conflict, there is no way to confidently assign it one label or the other. Nonetheless, considering an individual case may still be informative. Let us briefly consider the (at the time of writing) current war between Russia and Ukraine.

The government of Russia is reported to be offering soldiers salaries that are more than two times greater than the average Russian salary before the war (Stepanenko et al., 2022). Given that this compensation policy is enormously costly, it seems reasonable to assume that Russian leadership believes financial inducements play an important role in motivating at least some of its citizens to fight. If financial inducements became a less reliable tool, due to greater prosperity, then perhaps this would have led to a higher level of caution in deciding to initiate the war.

On the other hand, nationalism and other sources of "bonus utility" have clearly played a major role in driving support for the war, including among the majority of Russia's

population who will end up poorer as a result of the war. If this portion of Russia's population became less sensitive to economic damage, due to greater prosperity, then perhaps the country's leadership would have been even more comfortable in its decision to initiate the war. It may, for example, have worried even less about popular uprisings.

Ultimately, in my view, the most reasonable guess is that the war is scarcity-driven. The cooperation of combatants seems especially critical to the success of the war and the regime's stability. (When unsuccessful wars help to trigger regime changes, as in the case of the Russian Revolution, disaffected combatants of course tend to play especially large roles.) The financial inducements offered to combatants are also especially large and costly to the government, suggesting that the government believes they play an important role in ensuring cooperation. However, of course, this is an exercise in speculation.

5.4.2.2.3 What is prosperity's net effect? So far, in my view, rising prosperity has probably had an overall pacifying effect. I see two reasons for thinking this.

First, sizeable financial inducements for combatants are a fairly commonplace feature of interstate war. They also seem to have played an even more prominent role in preindustrial wars, at least in some regions and historical periods. For instance, in a discussion of medieval European warfare, Postam (1964) writes: "Military operations in the middle ages were expected to yield profits.... In certain periods, especially in years of active and successful campaigning, private profits of war could be large and enable fortunate individuals to enrich themselves; in all periods they figured very prominently as inducements to soldiers and as fuel for warlike sentiments." The prevalence of financial inducements at least weakly suggests that scarcity-enabled wars have tended to be more historically prevalent than scarcity-constrained wars.

Second, our bargaining model suggests that prosperity may have a *larger effect* on the risk of scarcity-driven wars than it does on scarcity-driven wars. Our quantitative analysis in Section 5.4.1.4 suggested that, at least when holding fixed certain parameter values, the

bargaining range tends to change more quickly for scarcity-enabled wars than for scarcity-driven wars.

To sum up: scarcity-driven wars appear to be both (a) more historically prevalent and (b) more strongly affected by changes in prosperity. Therefore, we should probably expect prosperity to have had a *net* pacifying effect so far.

However, this reasoning is still relatively tenuous and leaves room for reasonable debate. Furthermore, even if prosperity has had a net pacifying effect over the past few centuries, this does not guarantee that it will also have a net pacifying effect over the coming centuries. Since prosperity reduces the risk of scarcity-driven wars, but increases the risk of scarcity-constrained wars, it follows that—all else equal—scarcity-constrained wars ought to increasingly predominate as the world grows wealthier. Beyond a certain point, therefore, scarcity-constrained wars could become prevalent enough for the net effect of prosperity to reverse.

5.4.2.3 What about militarized interstate disputes?

When we explored anticipated effect sizes, in Section 3.4, we found that a conflict needs to have a fairly high expected mortality rate for the effect of prosperity to be large. Of course, given the range of simplifying assumptions the model makes, the particular numbers it outputs should not be taken too seriously.

However, this does raise the question of how we should expect rising prosperity to influence the prevalence of low-mortality interstate disputes. Most "militarized interstate disputes" (MIDs), as logged in the Correlates of War dataset, do not involve any fatalities. The disputes that do involve fatalities also typically involve only a small number of fatalities. One common example of a zero-fatality MID would be an intentional decision to briefly fly a military aircraft through another state's airspace.

Wars of course matter much more than these sorts of low-intensity disputes. However, since interstate wars are so uncommon, scholars often study the overall rate of MIDs as a

proxy for the likelihood of interstate war. Therefore, to empirically test this paper's theory, it is valuable to also understand its implications for low-fatality MIDs.

Naively applying our bargaining model to low-fatality MIDs suggests that the effect should be very small. The mortality risk is very low, and, most likely, both the financial and non-financial implications are small too.

However, changes in the risk of war could also have *flow-through effects* on the risk of low-fatality disputes. On the one hand, if states perceive the risk of war to be low, they may feel more comfortable taking provocative actions that qualify as MIDs.[13]

On the other hand, if states perceive the risk of war to be low, this may also reduce incentives or opportunities to initiate MIDs. Many MIDs — such as flying a military aircraft through another state's airspace — are plausibly explained by states attempting to signal they are unafraid of war, in order to deter their adversaries or compel them to accept unfavorable arrangements rather than fighting. If two states are actively guarding shared borders, policing key waterways, or monitoring each other's territory, this also simply increases the potential for MIDs to occur. If people in one state see another state as threatening, or are keen to go to war with that state, then military or political leaders may also find that initiating or condoning low-level disputes is a good strategy for winning domestic support.

Ultimately, the idea of a tight positive link between the risk of war and the rate of MIDs is is fairly intuitive. If we consider US-Canada relations and PRC-Taiwan relations, for example, it simply makes a great deal of sense that the pair of states with the higher risk of war also has a much higher rate of MIDs.

Overall, insofar as prosperity substantially reduces the risk of war, we should probably also expect it to substantially reduce the risk of MIDs in general.

[13]This kind dynamic is sometimes characterized as a "stability-instability paradox," borrowing a term that originated in the study of nuclear deterrence (Lindsay, 2016).

5.4.2.4 Summary

In this section, we have discussed a number of complications that were not represented within our simple bargaining model. First, we outlined more plausible and detailed causal pathways by which prosperity can affect leaders' decisions to go to war. Second, in light of the models' implication that prosperity can make certain categories of war more likely, we discussed the *net effect* of prosperity on the risk of war. Third, we more carefully examined the implications of prosperity for low-fatality disputes.

Overall, we have support for the hypothesis that rising prosperity has substantially lowered the overall risk of interstate conflict.

5.5 Other Effects of Growth

Of course, when states undergo economic growth, they do not simply become more prosperous. They also experience a wide range of more *structural* economic changes, such changes in their sources of wealth, trade relationships, and internal markets. As many scholars have noted, these structural changes may also influence the risk of war.

The purpose of this section is to briefly discuss several of these additional causal pathways linking economic growth to the risk of war.[14]

5.5.1 Changes in the Mobility and Fragility of Wealth

As states undergo economic development, the nature of their wealth tends to evolve. In the past, most states' wealth primary took the form of arable land, control of key points along trade routes, and access to a supply of agricultural labor. These forms of wealth are

[14]My focus throughout this section is on how the risk of conflict, for any given pair of states, might change after *they* experience growth. However, we could also consider the implications of *other states* growing. Some scholars, such as Gartzke and Weisiger (2014), have suggested that there could be important "spillovers" on the risk of conflict between third parties. One possible story is that economically developed states have a particularly strong interest in avoiding disruptions to global markets, so will be particularly inclined to put effort into preventing conflicts elsewhere in the world.

not very mobile and, with the partial exception of the agricultural labor supply, are not very easily damaged in the course of wars.

As a number of IR scholars have noted, the wealth of highly developed countries often has a very different character (Angell, 1913; Rosecrance, 1986; Brooks, 1999). The additional wealth that growth creates is often embodied in physical and financial capital and other more intangible forms of capital, such as the reputations, cultures, and business relationships of large firms. These forms of wealth are often either more fragile or more easily moved across borders. While an oil field might be seized in a ground invasion, without much loss, seizing McKinsey (without doing too much damage to its market cap) is a more problematic endeavor.

In the context of the bargaining model, this increase in the mobility and fragility of wealth manifests as a decrease in α and β parameters. A smaller portion of a highly developed state's wealth – although, presumably, still a larger total quantity of wealth – can be seized through conquest.[15]

5.5.2 Changes in Economic Interdependence

As states undergo economic growth, their economies typically become increasingly intertwined with those of other states. Trade volumes rise and specialization increases. These changes are mostly causes of growth, but they are also partly consequences of growth. For example, if some state's population mainly consists of farmers who produce just enough food to survive, this state is not going to be in a position to export very much of the output it produces or to shift a large portion of its population into some specialized industry. Growth also enables the development of transportation infrastructure that is necessary for high-

[15]In line with an observation made by Coe (2015), the higher *rates* of economic growth that have been characteristic of developed economies (at least until fairly recently) may also have increased the portion of wealth that two states can expect to destroy in a war. If a war actually damages a country's capacity to grow, rather than temporarily knocking it back along the same growth path, then this damage will compound over multiple years. For example, if conflict reduces a state's annual growth rate from 3% to 2%, then in ten years' time the GDP of this state will be a full 10% lower than it otherwise would be. In concrete terms, the market value of, for example, Apple largely reflects expectations of future revenue growth; destroying its capacity for growth would mean destroying most of its value.

volume trade.

It has long been argued that growing interdependence can be expected to mitigate the risk of war (Angell, 1913; Gartzke, 2007; Coe, 2015). The most common version of the argument notes that, if two states are economically intertwined, then neither state can damage the other's economy without damaging its own. Suppose, for instance, that car manufacturers in State A import essential parts from State B. If State A destroys the factories that manufacture these parts, in the course of a war with State B, then it will also have damaged its own auto industry. The damage spreads and compounds in a way it would not if the two economies were not intertwined. As Angell (1913) memorably puts this point: "If in the time of the Danes, England could by some magic have killed all foreigners, she would presumably have been better off. If she could do the same thing today, half her population would starve."[16]

In the context of our bargaining model, then, an increase in economic interdependence manifests itself as a decrease in the α and β parameters. A larger fraction of the wealth that exists at the start of any war can be expected to be destroyed.

5.5.3 Changes in the Prevalence of Democracy

Somewhat controversially, many scholars believe that economic growth tends to support democratization. Since democratic states are much less likely to fight each other, this hypothesized effect — if real — would imply another possible pathway by which growth reduces the risk of conflict.

As summarized in a recent review paper by Treisman (2020), there has historically been a strong cross-national correlation between a country's GDP-per-capita and its level of democracy (e.g. as measured by the Polity2 scale).[17] Although this claim is more con-

[16]However, as Gartzke and Zhang (2015) point out, this argument seems to rely on the level of economic dependence and the level of economic disruption both being very large. Since war is already enormously costly, to start, "[m]arginal increases in the overall cost of fighting can have at most a marginal effect on whether conflict occurs."

[17]From 1850 to 1972, the correlation within individual years ranged from a minimum of $r = .39$ to

troversial, time series data also appears to at least weakly suggests that this correlation is not just a product of certain fixed traits countries (e.g. geography) that contribute to both development and democratization.[18]

Longer-run historical patterns are also seemingly consistent with the hypothesis that economic development supports democratization. It is notable that, although democratic city-states or smaller communities have been common in different places and times, large-scale states were virtually always autocracies up until the past few hundred years. The emergence of modern economic growth has lined up relatively closely with the emergence of modern democracy.

Furthermore, a number of authors have put forward theoretical mechanisms that could explain why it is easier for economically developed countries to democratize (e.g. Boix, 2003; Acemoglu and Robinson, 2005). For instance, in pre-industrialized countries, rent earned from owning agrarian land (and having non-elites work it or pay for the privilege of working it) is a much greater source of the ruling elite's wealth (Crone, 2015). This may lead to a dynamic where political elites have a particularly large amount to fear from democratization, since non-elites have a natural tendency to call for land reform policies that redistribute ownership of land to themselves. In an industrialized country, on the other hand, the calls for the redistribution of the relevant forms of wealth may be less intense for a variety of reasons. One potential reason is that heavily redistributing ownership of companies or capital, or setting very high tax rates on income earned from capital investments and labor, would tend to substantially reduce investment and economic growth and leave the country poorer overall. Other aspects of or partial downstream effects of economic development that have been raised as potentially important include urbanization (e.g. because it increases

a maximum of $r = .74$. It declined substantially post-1972, although this decline seems to have been substantially driven by the emergence of oil economies that exhibit high GDP-per-capita without the other trappings of economic development.

[18]Treisman reports that — although a country's current GDP-per-capita is not clearly predictive of the chance it will democratize within the next year — it is predictive of whether it will democratize within the next one or two decades. He suggests that this is consistent with economic development having an effect on the chance of democratization that takes some time to materialize.

the extent to which political elites need to fear revolutions), the emergence of new wealthy groups outside the political elite (e.g. because these groups have incentives to constrain the powers of political elites who might wish to appropriate their wealth), and rising education and the growing need for an educated workforce (e.g. because this increases the potential for informed political participation by non-elites and also gives elites less reason to object to state support for non-elites).

Ultimately, the relationship between economic development and democratization is murky. Development is at most one factor that contributes to democratization. There may also in part be a causal relationship running in the opposite direction, with democratization supporting growth (Acemoglu et al., 2019). It is also entirely plausible that there are common factors — for instance, cultural and intellectual factors — that help to explain the co-emergence of both modern democracy and modern growth.

That being said, it ultimately seems likely that economic development supports democratization to at least some degree. Since it is widely believed that democratization makes pairs of states less likely to fight each other, an effect on the prevalence of democracy can in turn imply an effect on the prevalence of war.[19]

One possible way we can understand the effect of democratization, in the context of the model, would be as a shift in how idealized the model's assumptions are. The model assumes that states are perfectly responsive to the interests of the people they contain. As states democratize, this assumption may become less badly wrong; the mortality risks and non-material costs of conflict may be more fully factored into states' aims, in a way that tends to expand the "bargaining range" and make conflict less likely.

The effect of democratization could also be directly modeled as a reduction in the values of the B_i terms for a pair of states, if we assume that the reason democratic states fight less at least partly has to do with ideology or other non-material preferences. For example, it

[19]One caveat here, however, is that there are some indications that one-sided democratization — where only a single state democratizes, while another remains authoritarian — may in fact increase the risk of conflict within the dyad. Gartzke and Weisiger (2014) comment on the potentially mixed effects of rising democracy on the global rate of conflict.

may be the case that people in the United States would be less comfortable fighting Iceland than they would be fighting Iran *in part* because Iceland is also a democracy.

Barnhart et al. (2020) have also suggested that democratic states may be less likely to fight each other because — once they give women suffrage — they are more responsive to the preferences of the average women. (Political elites are, of course, normally men.) Women tend to be more averse to violence and war than men are, implying lower values for B_i terms than one would see in an all-male society. Female suffrage could therefore be especially relevant for driving democratic peace.

5.5.4 Changes in Culture

Many anthropologists and historians have argued that economic growth tends to bring changes in cultural values (Morris, 2017; Inglehart and Welzel, 2005). Reporting on the results of the long-running World Values Survey, Inglehart and Welzel write: "Four waves of survey data from more than eighty societies demonstrate that socioeconomic development tends to transform people's basic values and beliefs – and it does so in a roughly predictable fashion." They summarize these transformations as a shift toward "secular-rational" values, away from "traditional values," and toward "self-expression" values away from "survival" values.

Gat (2017), drawing on this research, suggests that the appeal of war declines as individuals become more individualistic, tolerant, and pre-occupied with self-fulfillment. Mousseau (2009) presents a similar argument that individuals in contract-intensive economies "learn to respect the choices of others and value equal application of the law," with one result being increased respect for sovereignty and international law. Braumoeller (2019) notes, however, that not all recent shifts in values necessarily reduce the risk of war. For instance, concern for human rights and individual freedom can motivate military interventions that would not otherwise be in a state's interest.

Changes in values can be modeled as increases or decreases in the B_i parameters. If

people develop a strong distaste for violence, for example, then they will tend to assign lower non-material value to war.[20]

5.5.5 Changes in Young Men's Access to Casual Sex

Rape is nearly always prevalent in war, even though it often less discussed than other forms of violence (Leatherman, 2011). For example, during the Second World War, victorious Soviet soldiers committed mass rape while they occupied Germany; the number of women raped is sometimes estimated to be as high as two million (Naimark, 1995). Another example from the Second World War is the Japanese military's sexual enslavement of tens or hundreds of thousands of largely-Korean "comfort women," to serve Japanese troops. Although these cases are particularly infamous, they are not, unfortunately, anomalous. Rape has always been a part of conflict, and, in premodern conflict, tended to be even less veiled (Gat, 2008).

In addition to offering men opportunities to commit rape, conflicts have also offered them greater opportunities to have sex with prostitutes. For instance, during the Vietnam War — which also contained many documented cases of rape by all sides — an enormous sex industry arose to serve American servicemen (Cottrell, 2015). This industry was enabled by the economic desperation of many women and families affected by the war, by at least tacit endorsement by military leaders concerned with maintaining servicemen's morale, and, likely, by the fact that the servicemen (temporarily deployed abroad and placed in a community of other young men) were less at risk of suffering reputational damage from their sexual behavior than they would have been at home.

It seems likely, overall, that opportunities for sex have played at least *some* role in incentivizing some young men's initial or continued participation in war.[21] Grimly, we can

[20] One general concern about proposed cultural pathways from economic development to peace is that they confuse cause and effect. The modern transition from economic stagnation to growth has often been partly attributed to cultural change (Clark, 2008; Mokyr, 2016). A separate concern is that a society's attitude toward violence may largely reflect that society's material incentives to use violence, rather than having much independent explanatory power (Morris, 2017). One imagines, for instance, that elites in medieval Europe might have placed less cultural emphasis on fighting if fighting was never profitable for them.

[21] Rape is also, in some cases, a tool of war. For example, the implicit threat of rape can be used to terrorize

also represent this incentive in our model as a contributor to the B_i term. The larger the gap is between men's sexual opportunities when at war and their sexual opportunities at home, the larger this additional sexual incentive for conflict will tend to be.

As Gat (2017) has argued, if a country's sexual norms become looser — such that women become more free to participate in casual sex without needing to fear serious personal consequences — then men's sexual incentives to participate in conflict seemingly should decline. The more available casual sex is to men by default, the less their sexual opportunities will tend to increase — or, beyond a certain point, the more their sexual opportunities will actually tend to *decrease* — if they leave their homes to participate in a conflict.

On average, countries appear to experience a loosening of sexual norms as they undergo economic development, among the other cultural changes they tend to experience. For example, a country's GDP-per-capita strongly correlates with its population's attitude toward premarital sex (Jung, 2016). Of course, as well, sexual norms in most economically developed countries are looser than they were before they became developed.

If the relationship between development and sexual liberalism is in fact causal, then this suggests that economic development may reduce sexual incentives for conflict. This would manifest, in our model, as a decline in the B_i terms.[22]

5.5.6 Changes in Demographic Structure

National fertility rates (i.e. the average number of children a woman will have) tend to decline as countries grow wealthier (Kirk, 1996). This pattern is nearly universal and

a population, and gang rape may serve the purpose of building trust among male combatants, similar to the social purpose of initiation rituals in criminal gangs (Cohen, 2018). Rape may also in some cases be an expression of sadism as much as it is an expression of sexual desire. Even so, however, it seems likely that people who choose to commit rape tend to derive pleasure from rape (Pinker, 2002).

[22]It also seems plausible that, as countries become more economically developed, negative attitudes toward rape become increasingly strong. One slightly suggestive finding, for example, is that a country's level of economic development explains roughly 30% of cross-country variance in the rate of self-reported intimate partner violence (Heise and Kostadam, 2015). If a country's population becomes less tolerant of rape as it undergoes economic development — such that men are more uncomfortable with committing rape or have greater reason to fear social stigma from committing rape — then this would also reduce the sexual incentives for conflict.

results in stark cross-country differences. The highest fertility rates in the world are all in low-income countries such as Niger (6.7) and Chad (6.1), while the lowest fertility rates in the world are strong disproportionately in high-income countries such as South Korea (.9) and Singapore (1.0) (United Nations Population Fund, 2023). Although there is no consensus explanation for why fertility rates tend to decline as countries grow, the relationship is widely thought to be causal.

At the same time, of course, people in wealthier countries tend to have greater life expectancies, due to better medical care, public health, and nutrition.

The result of declining fertility and rising life expectancy is that young people — most consequentially, young men — tend to make up a smaller portion of a country's population as it grows.

Of course, as evidenced by violent crime rates, young men are especially inclined toward violence. There is also some specific evidence from public opinion surveys that they tend to be especially supportive of war (e.g. Rosensteil, 2006). For this reason, Gat (2017) has suggested that, as young men become less common within a country, that country will tend to become less disposed toward war. This effect can be represented in our model as a decline in the values of the B_i terms.

5.5.7 Changes in the Offense-Defense Balance

Economic growth may also influence the "offense-defense balance," meaning the relative ease of attack and defense (Jervis, 1978), through a number of different pathways.

First, as Gartzke (2007) has noted, growth may reduce the degree of protection provided by geographic obstacles to invasion. Two states with modern naval technology and massive military budgets can plausibly fight a competitive war even if they are separated by an ocean. Two states with pre-modern naval technology and miniscule military budgets cannot. More generally, growth supports higher military spending, and increasing military spending from very low levels will often initially tilt the offense-defense balance toward of-

fense (Garfinkel and Dafoe, 2019).

Second, growth alters the terrain that a ground war would be fought on. Developed countries typically have higher levels of urbanization, lower levels of wilderness, and better transportation infrastructure. Although these changes can cut both ways, the difficulty and danger of offensive operations in urban environments is sufficiently large to suggest that defenders benefit on net (Rosenau, 1997).

Since our bargaining model does not distinguish between the attacking state and the defending state, there is no direct way to represent changes in the offense-defense balance within the model. However, if we assume that the "victor" in the war manages to seize assets from the other state, then a shift toward defense-dominance should manifest as a decrease in α (the portion of initial wealth the victor will hold at the end of the conflict). This is because, as defense dominance increases, any attempt to seize wealth from the other state must become either be less successful or more costly.

5.5.8 Changes in the Information States Have About Each Other

As noted above, in the context of bargaining theory, "information asymmetries" are generally thought to increase the risk of conflict. The basic reason for this is that, when states do not have the same information, they can end up with different beliefs about where the "bargaining range" lies. For example, if two states both believe that the other state will probably lose badly if they fight, then anything that looks like a reasonable bargain to either state may look unreasonable to the other state. The result is that it may be very difficult for them to reach a peaceful compromise that they both recognize as preferable to war.

In principle, therefore, economic growth could also reduce the risk of conflict by reducing information asymmetries. There in fact seem to be some pathways by which economic growth could at least plausibly have this effect. For example, Kaplow and Gartzke (2021) have suggested that it may be easier to assess the military capabilities of economically developed states from the outside, in part because these capabilities are more likely to flow

from the possession of expensive weapons platforms that cannot easily be kept secret. Pairs of developed states might also know more about each other because trade relationships and travel cause relevant people to interact more heavily; because information technology adoption and the rise of mass media makes countries more transparent to outsiders; or because rising government budgets and new technologies tends to facilitate intelligence-gathering capabilities.

Gartzke et al. (2001) have further suggested that a state's participation in international capital markets (which is in part facilitated by growth) could make it easier for other states to discern if it is *bluffing* when it suggests that it would be willing to go to war. The argument here is that investors are less willing to invest in states that they expect to go to war; any state that signals its willingness to go to war, therefore, is doing this with the knowledge that it must suffer a serious financial cost simply to send the signal; the fact that the state is willing to suffer this cost then makes its signal more *credible*. Ultimately: the more easily states can communicate their actual intentions and expectations without the other state suspecting them of bluffing, thanks to capital markets, the more easily they can identify genuinely mutually beneficial arrangements. Others scholars (e.g. Dafoe and Kelsey, 2014) have suggested, however, that the transparency implications of capital markets are unlikely to have a *large* effect on the risk of conflict.

5.5.9 Should We Expect the Marginal Effect of Growth Decline?

When we explored the Diminishing Returns Mechanism, above, we noted that it suggested that the marginal effect of growth might decline as states become wealthier. One basic reason to expect this was that "scarcity-driven wars" may become less common as the world grows wealthier.

A number of the additional mechanisms discussed in this section also suggests that the marginal effect of growth will at least eventually decline. For example, within an individual dyad, variables such as the level of democracy, the rareness of young men, the availability

of casual sex, and the portion of wealth that is difficult to steal presumably cannot increase infinitely. Although some of these mediating variables may be able to increase for quite a while longer, ceilings are presumably going to be approached at *some point*.

This suggests that even if growth has *historically* been an important pacifying force, it may at least gradually become less relevant to the risk of conflict.

5.5.10 Summary

When countries experience economic growth, they do not simply become more prosperous. They also experience a range of structural and qualitative changes that may influence the risk of war. Some of these changes could conceivably exacerbate the risk of war. However, for the most part, they also seem to push in the direction of greater peace.

At the same time, it is worth keeping in mind: even if these additional causal mechanisms have historically supported peace, they may become less important as the world becomes wealthier. We should not be surprised if the marginal effect of growth is smaller in an already-wealthy world.

5.6 Growth and the Correlates of Interstate War

5.6.1 Previous Studies

In sharp contrast to the civil war literature, most relevant papers in the interstate war literature report that there is no robust association between GDP-per-capita and conflict (Maoz and Russett, 1992; Mousseau and Hegre, 2003; Gartzke, 2007). The relationship is typically reported to become statistically insignificant once basic controls are introduced.[23]

[23]Bremer (1992) reports an early contrary result. He includes a binary variable "both more economically advanced" in a multivariate Poisson regression and finds a significant relationship with interstate peace ($p = .005$). However, this binary variable does not strongly correspond to prosperity. It is developed by calculating an "economic capability" index based on state's energy consumption and iron and steel production, calculating a "demographic capability" index based on a state's urban population and total population, and then classifying a state as "economically advanced" if its share of the world's total economic capability in a given year is greater than its share of the world's total demographic capability. One issue with

Researchers do, however, *typically* report fairly robust and statistically significant relationship between trade and peaces (Gartzke and Zhang, 2015). For this reason, insofar as the causes of war literature has explored the consequences of economic changes that are (at least in part) downstream of economic growth, it has focused most heavily on expansions in trade.

Other hypothesized causal mechanisms relating growth to interstate peace are typically treated as having more tenuous empirical support. Often, proponents of these theories rest their arguments on interaction effects between GDP-per-capita and other variables. For instance, Mousseau et al. (2003) observe an interaction effect between GDP-per-capita and democracy to argue for his culture-based theory of peace. Gartzke (2007) has also highlighted an apparent interaction effect between GDP-per-capita and contiguity: he reports that there is a negative relationship between GDP-per-capita and conflict for contiguous states, but a positive relationship for non-contiguous states. He hypothesizes that development reduces the economic value of conquest, while also making it easier for distant states to fight.[24] Finally, a non-linear relationship between GDP-per-capita and conflict has also been proposed, with Boehmer and Sobek (2005) arguing that the relationship does ultimately become negative but only above a certain development level.

Overall, unlike in the civil war literature, there is no general academic consensus that — other than through its positive relationship with trade — prosperity is actually meaningfully associated with interstate peace. Papers tend to either report that there is no robust association or rest their cases on somewhat tenuous interaction effects. In a relatively recent review article, Schneider (2017) simply states, "[D]yadic development does not reduce the probability of interstate war."

One simple indication of the field's overall skepticism that growth matters is that authors seldom decide to include GDP-per-capita as a control in papers studying other

this approach, for example, is that it implies that the portion of states that are "economically advanced" has not meaningfully increased over time.

[24]However, the robustness of this interaction effect has been questioned (Choi, 2011).

possible causes of conflict. For instance, in a recent landmark paper on whether female suffrage reduces the risk of interstate war, Barnhart et al. (2020) include more than ten controls for other variables that are widely seen as predictive of war — including controls for trade, democracy, nuclear weapon position, and differences in military capability levels — but do not include one for GDP-per-capita.

5.6.2 The Importance of Fixed Effects

One crucial point about the empirical literature on economic development and interstate war is that virtually none of the relevant papers present *fixed effect models*. Wealthy states differ from poor states along an enormous number of dimensions. Many of these differences are "fixed," in the sense that they were not caused by growth and have been fairly stable over time. Economists have identified dozens of fairly stable geographic, institutional, and cultural factors that are predictive of a country's level of development (Weil, 2009). Examples of such factors include whether a country is landlocked, whether a country has inclusive political institutions, whether a country has large natural resource endowments, whether a country has a high-trust culture, and whether a country has a history of colonialism. Although it is quite plausible that some of these factors also influence the risk of interstate war, or closely correlate with other factors that do, we cannot hope to introduce well-measured individual controls for *all* of them. It is not plausible that including any small set of control variables will be sufficient to fully screen off all of the relevant fixed differences between pairs of states.

Dyad-level fixed effects models solve this problem by introducing a dummy variable for each pair of states (Green et al., 2001). This allows us to control for any fixed properties of pairs of states that influence both their average national income levels and their propensities to enter into conflicts. For example, this design choice ought to nearly eliminate confounding from geographic properties of states (such as the property of being landlocked) or from cultural traits that persist for many decades. Ultimately, including fixed effects allows us

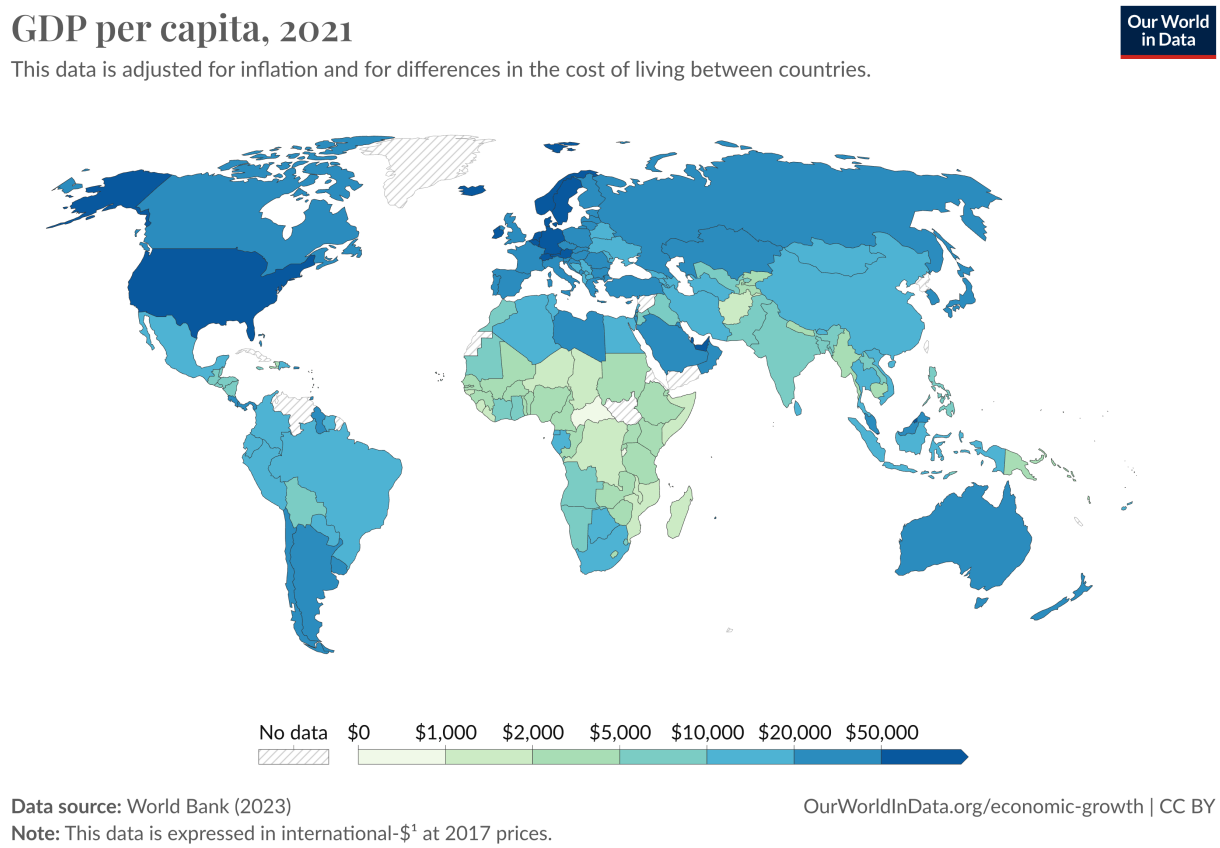

Figure 5.5: A map of the world, showing variation between the income levels of different states in 2021 (Roser et al., 2023). It is clear, when one reflects on the patterns on this map, that wealthy and poor states tend to differ along many geographic, historical, institutional, cultural, and international-status-related dimensions. This suggests that adding a small number of control variables, such as democracy scores, is unlikely to come close to removing omitted variable bias.

to focus on the question of whether pairs of states become less likely to fight as they grow wealthier.

The field of international relations has, historically, had a substantial degree of aversion to the use of fixed effect models, although this aversion has declined over time (Erikson, 2009).[25] Since these models essentially "screen off" all information regarding fixed differences between pairs of states, they tend to report larger standard errors and are therefore

[25]To give an example, one widely-cited 2010 paper by leading scholars of international conflict — in this case, analyzing the effect of trade on the risk of conflict — offers this brief comment on the use of dyadic fixed effects: "Dyadic fixed effects are not appropriate because...all dyads that never fought would be dropped, and all between-dyads variation would be ignored" (Hegre et al., 2010).

more likely to report null results even when genuine causal relationships exist. For example, suppose that having a mountainous border lowers the risk of war. A fixed effect model will have trouble detecting this relationship, since whether or not a pair of states has a mountainous border is (in all but a few cases) a fixed trait; we simply do not have enough relevant cases to tell whether states become less likely to fight when their borders shift so as to become mountainous. Likewise, when we use fixed effects models, we cannot glean any information from pairs of states that have never experienced a conflict, since there is simply nothing that indicates whether these pair of states have become more or less conflict-prone over time; as a result, all pairs of states that never experience a conflict are dropped from our sample. The fact that fixed effect models seem to "throw away information" in these ways has historically been used as a justification for avoiding them (Beck and Katz, 2001; Oneal and Russett, 2001). For instance, in one classic article published in the top International Relations journal, Beck and Katz (2001) call the use of fixed effects in models of interstate conflict "clearly silly," "clearly a bad idea," "never ideal," "hardly...very interesting," "absurd," "simply nonsense," reflective of a "nihilist position," and "*never* a good idea."

To be clear: the central problem here is that fixed effect models are more likely to produce *false negatives*, by failing to detect some genuine causal relationships. Importantly, though, their corresponding virtue is that they are also (in typical cases) less likely to produce *false positives*. The fact that fixed effect models "screen off" information about fixed differences is not, in itself, a reason to worry that their coefficient estimates for independent variables will be *biased* (Cook, 2020).[26] On the other hand, we have very good reasons

[26]One way that bias *could* arise is if growth effects dyads differently, depending on whether or not the dyad has a substantial chance of experiencing a conflict. Since dyadic fixed effect models cannot be used to glean any information from dyads that have never experienced a conflict, we are only learning about the effect of growth on dyads that are on average at least moderately conflict-prone. However, in practice — even if growth *does* affect these two kinds of dyads differently — this is not much of a problem, because we primarily *care* about the effect of growth on dyads that have a meaningful risk of fighting. If growth actually increases the annual chance of an unprecedented conflict between Singapore and Iceland from 1-in-1,000,000 to 1-in-100,000, for example, then this shift would not be of much practical importance. Dyadic models would still be telling us about the effect of growth on the sorts of dyads we are concerned about.

to expect fixed effect models to correct for sources of omitted variable bias that alternative models do not. Therefore, even if we do not *exclusively* rely on fixed effect models, we should typically at least consider them and compare them to the results of pooled models.

One statistics textbook explains the situation in the following way, in the context of studies of the effect of income on individual behavior:

> *Fixed effects methods completely ignore the between-person variation and focus only on the within-person variation. Unfortunately, discarding the between-person variation can yield standard errors that are considerably higher than those produced by methods that utilize both within- and between-person variation. So why do it? The answer is that the between-person variation is very likely to be contaminated by unmeasured personal characteristics that are correlated with income. By restricting ourselves to the within-person variation, we eliminate that contamination and are much more likely to get unbiased estimates. So what we're dealing with is a trade-off between bias and sampling variability. For nonexperimental data, fixed effects methods tend to reduce bias at the expense of greater sampling variability* (Allison, 2009).

One practical upshot is: If fixed effect models report that pairs of states become less likely to fight as they become wealthier, then — because estimates produced by fixed effect models *tend* to be less biased — this finding should overshadow the null results commonly produced by pooled models.[27]

Surprisingly, however, I am not aware of any major papers that examine the relationship between GDP-per-capita and interstate peace using fixed effect models at the dyadic level.[28][29] In this section, I attempt to partly address this gap in the literature by running

[27]On the other hand: since fixed the larger standard errors associated with fixed effect models make them more prone to false negatives, null results produced by fixed effect models should not necessarily overshadow significant results produced by pooled models.

[28]Gartzke (2007) is notable for including regional dummy variables, for instance a control variable that indicates whether a dyad is located within Europe, but does not include dyad-level controls. It is also unclear whether or not models that include regional dummy variables should be expected to produce less biased estimates than models that exclude them (Dafoe, 2011)

[29]A priori, it seems unlikely that no major paper would have reported the relevant regression. However, I have not been able to find any paper that does this. Tellingly, even other relatively recent works that argue that growth may matter — which would have strong reasons to report positive results from regressions with fixed effects, if they were aware of these results — tend to only report the results of regressions run on models without fixed effects. For example, Gat (2017) argues at length that growth (technically, "industrialization") reduces the risk of interstate conflict; so far as I can see, though, none of the papers he cites estimate the effect of GDP-per-capita using dyad-level dummy variables. Similarly, Gartzke and Weiseger (2014) are

regressions used fixed effect models.

5.6.3 Research Design

Hypothesis

I will use the models to test the hypothesis that *pairs of states become less likely to fight as they become wealthier*.

Insofar as there is a positive association between economic development and interstate peace, I will also test the hypothesis that *this association is not wholly or primarily the result of growth's association with trade and democracy*. Trade and (joint) democracy are already widely believed to have pacifying effects. Even if growth actually does cause peace, therefore, this causal relationship would be less interesting if it was entirely mediated by growth's effect on trade and democracy.

Data Sources

For my conflict data, I will draw from Version 5.0 of the Correlates of War Militarized Interstate Dispute Dataset (Palmer et al., 2022). For my GDP-per-capita data, I will draw from the 2020 version of the Maddison Project Database (Bolt and van Zanden, 2020). I will slightly modify the Maddison Project Database, however, by filling in gaps in the GDP estimates for individual states by assuming exponential growth within these gaps.[30]

sympathetic to the hypothesis that growth supports peace, but none of the papers they include in their literature review contain this estimate either. The first ten pages of Google Scholar results for papers that include the phrases "GDP", "Correlates of War", "fixed effects", and "interstate" also do not turn up any examples of papers that run the relevant regression. One possible explanation for this omission that I can see is that: the result was initially not published because reliable long-run GDP-per-capita data was unavailable, with the first widely-used dataset in International Relations appearing only in the early 2000s (Gleditsch, 2002); then was not published because the use of fixed effects in International Relations was rare throughout the mid-2000s; and then was not published because the relevant literature became somewhat dormant, because it was now assumed that growth had no effect on dyads, or because the result was too basic to warrant publication.

[30]Assuming exponential growth within these gaps would be highly problematic if we were studying the effect of *growth rates* on conflict. However, my expectation is that filling in the gaps in this way will lead to more accurate results than we would produce by simply dropping dyad-years with missing GDP-per-capita estimates from the dataset. Since log(GDP-per-capita) — which we will be using as our independent variable — does not change much from year to year, and since growth tends to be relatively steady in any case, these

Importantly, the Maddison Project Database GDP-per-capita estimates should not be taken as gospel. Since pre-WWII economic data is scarce or unreliable for many countries, a fair amount of speculation inevitably goes into the production of all longrun GDP-per-capita estimates. Nonetheless, particularly since we will be log-transforming the estimates, I expect that the estimates are still *reliable enough* to be useful. If nothing else, it would be surprising if the estimates were not at least very highly correlated with the true historical values of countries income levels.

The data for the control variables in the regressions are drawn from a number of standard supplementary datasets. Since I have used the peacesciencer package to construct the overall dataframe, the reader can consult the documentation for this package for a list of these datasets (Miller, 2022).[31] Notably, since the supplementary datasets do not typically extend back to 1816, including controls tends to reduce the number of available datapoint and reduce the temporal scope of the analysis.

Dependent variable: involvement in militarized interstate disputes

We will use *War* as a dependent variable. *War* is coded 1 in the first year of dispute that ultimately rises to the level of a "war" involving the participants. Using criteria established by Correlates of War team, this means the dispute must ultimately involve sustained combat, involve organized forces, and result in at least one thousand battle-related fatalities within a twelve-month period.

Fortunately for the world, but unfortunately for our analysis, there are not very many wars. This often makes it difficult to detect statistically significant effects. For this reason, scholars of conflict most often use militarized interstate disputes (MIDs) as the dependent variable in statistical analysis. (There are far more MIDs than wars.) We will therefore also

interpolated estimates should still be approximately correct. Note, as well, that the main alternative GDP-per-capita dataset that has been used in International Relations also opts to interpolate missing datapoints (Gleditsch, 2002).

[31] The one exception is that I have constructed the measure of economic dependence — trade flow divided by GDP — by drawing on the Maddison Project's estimates for the denominator.

use *MID* as an alternative dependent variable.

The use of MIDs in this case primarily justified by the assumption that, if growth does reduce the risk of two states experiencing a war, then it will likely also reduce the risk of them experiencing a MID. See Section 5.4.2.3 for a discussion of this point.

Independent variable: GDP-per-capita

Our primary independent variable is *AvgLogGDPPC*, the average value of the natural log of GDP-per-capita for each dyad-year.[32] It is intended to serve as a measure of the average level of development within a dyad.

This variable is lagged by one year, to reduce ambiguity about the direction of any apparent causal relationship between development and conflict. Unless stated otherwise, all other time-dependent independent variables are also lagged by one year.

Fixed effects

To mitigate concerns about confounding, as discussed above, I will use a logistic regression model with a dummy variable for each pair of states.[33] This allows us to focus our attention on the question of whether pairs of states become less likely to fight as they grow wealthier.

In some models, I also include year-level dummy variables to control for temporal fixed effects. Since GDP-per-capita tends to rise over time, variation in the overall conflict-proneness of different time periods could produce spurious correlations between GDP-per-capita and conflict. For example, if some other trend is making the world more peaceful over time, then we might observe states become more peaceful as they become wealthier — even if growth has no effect on peace. We also might also observe spurious correlations

[32]Average wealth, rather than average annual income, might in principle be a more relevant independent variable to consider. However, there is a very close relationship between income and wealth, and available estimates of average national wealth throughout history are far less complete.

[33]This is estimated using the fixest package and feglm function in R (Berge et al., 2021).

if the international system is simply, by chance, more or less stable in a particular period. Controlling for temporal fixed effects helps to address this issue.

In most of the models, however, I instead use the continuous variable *Year* and the additional polynomial terms $Year^2$ and $Year^3$ as temporal controls. This alternative approach is recommended in Carter and Signorino (2010). I am comfortable primarily using this approach because, as will be seen below, it tends to produce results that are extremely similar to the temporal fixed effects approach.

Other controls

We also include, for some models, the following common controls, which are often found to be predictive of the risk of interstate conflict. To reduce my own degrees of research freedom, this selection of controls is based on the selection of controls used in Barnhart et al. (2020).

- *AvgDepend* is a measure of economic interdependence. Specifically, it equals average of the bilateral trade-to-GDP ratios for the pair of states.[34]

- *AvgInclusion* is a variable that records the average Polity scores of the two states, shifted so that they are on a scale from 0 to 20. These scores describe the extent to which states are democratic, rather than autocratic.

- *EitherNuclear* is a binary variable that indicates if either state possesses nuclear weapons. *JointNuclear* indicates if both states possess nuclear weapons.

- *Ally* is a binary variable that indicates if the dyad holds any type of alliance, as reported in the COW Military Alliance dataset.

[34]It is more common to use the minimum bilateral trade-to-GDP ratio as a control. However, given that the primary independent variable is an average quantity, I believe that it is more natural to use average quantities as controls. In any case, robustness checks show that this particular research design choice has a negligible impact.

- *InterestSimilarity* is a measure of similarity between the two states' policy portfolios, as estimated in Signorino (1999).

- *CapRatio* is a measure of the gap between the two states' military capabilities. Specifically, it equals the natural logarithm of the ratio of their Composite Index of National Capability (CINC) scores in the Correlates of War National Capabilities dataset.

- *MajorPower* is a binary variable that indicates whether either state in the dyad is a major power, as reported in the COW State System dataset.[35]

- *Contiguity* is a binary variable that indicates if the states share a border or are separated by sufficiently small distance over water.[36]

- *Distance* is a variable that measures the natural logarithm of the distance between the capital cities of the two states. I also include *Distance^2* and *Distance^3*.

- *Year* records the year. In cases where I use this control, I also include *Year^2* and *Year^3* as recommended in Carter and Signorino (2010).

- *Peace Years* counts the number of years since the last militarized dispute within the dyad, as coded in the COW dataset on interstate conflict. I also include *Peace Years^2* and *Peace Years^3*.

Some of these control variables may be partly downstream of economic growth. Most notably, as discussed above, economic interdependence is probably in part an effect of growth. Growth may also support democratization (Treisman, 2020). Including controls for interdependence and democracy therefore screens off two causal pathways by which growth might influence the risk of war. If national income still predicts peace, even once interdependence

[35]The vast majority of states have consistently been minor powers: only nine have had major power status at any point in the time period covered by the dataset.

[36]On rare occasions, changes in borders will create or remove contiguity. Otherwise, including dyadic fixed effects already effectively controls for contiguity.

is controlled for, then this would count as evidence that rising interdependence and democratization are not the only pathways through which growth produces peace.

It is also worth keeping in mind that including more controls will not *necessarily* lead to more accurate effect size estimates. Collider bias is an important concern. For instance, if we control for possession of major power status, then lower levels of national income may begin to indicate higher levels of military ambition: a poorer state must work harder to become a major power.

It is valuable, therefore, to run regressions with both smaller and larger sets of controls. One cannot easily say which models are more or less biased. Introducing additional controls also, of course, tends to reduce statistical power.

5.6.4 Results

First, Tables 5.2 and 5.3 each report the results of five regressions on models that *do not* include dyad-level fixed effects.[37] The first table reports results with *MID* as the dependent variable, and the second table reports resuls with *War* as the dependent variable. For most of the MID regressions, there is actually a statistically significant relationship between AvgLogGDPPC and the risk of a MID. However, as more controls are added, the sign of the coefficient switches. Once the full set of controls is added, it becomes statistically insignificant. Then, the relationship between AvgLogGDPPC and war becomes insignificant once any basic controls are added. Overall, in keeping with previous published work, the pooled regressions do not suggest a clear and robust relationship between AvgLogGDPPC and conflict.

Next, Tables 5.4 and 5.5 each report the results of five regressions on models that *do* include dyad-level fixed effects. These results strongly suggest a negative relationship between AvgLogGDPPC and conflict. For the MID regressions, the relationship is *consistently negative and highly statistically significant* for all sets of controls. For the war regressions,

[37] All standard errors reported in the regression tables in this paper are clustered by dyad.

the relationship is consistently negative, and it is statistically significant for all of the models except the one that uses the full set of controls. Even in this case, however, the estimated effect continues to be large, suggesting that it may simply be hiding behind large standard errors.

The estimated effect size is also fairly substantial. Focusing on the MID model that includes the full set of controls: **When a pair of states becomes twice as prosperous, the risk of them experiencing a conflict falls changes by approximately -27% [-41%,-11%]**. Focusing on the war model that includes all of the controls give us -38% [-67%,+19%].[38] Given that the typical pair of states has become about *four times* as prosperous over the past century, this estimated effect size could be practically consequential.

It is also worth noting that, in the MID case, introducing controls for democracy and trade only very slightly changes the estimated effect size. In the war case, there is actually a substantial (albeit still not statistically significant) change in the estimate, but two-thirds of the estimated effect still remains after the inclusion of these controls. This observation suggests that the negative relationship between growth and conflict is mostly *not* explained by growth's association with democratization and rising trade.

In summary, when a pair of states experiences economic growth, its probability of experiencing a conflict appears to decline substantially. Particularly in the case of militarized disputes — but also, to a large extent, in the more specific case of war — the relationship is highly robust to the inclusion of standard controls and is *not* primarily explained by rising economic interdependence or rising democracy.

5.6.5 Discussion

These above results are consistent with the hypothesis that economic growth substantially reduce the risk of interstate war. Of course, they do not *prove* that growth supports peace. It is possible that some other unobserved variable both causes states to grow and

[38] See the appendix of this chapter for an explanation of how these estimates were calculated.

Table 5.2: Logistic regression *without* dyadic fixed effects. MIDs.

	Dependent variable:				
	MID				
	(1)	(2)	(3)	(4)	(5)
AvgLogGDPPC	−0.119*	0.327****	0.320****	0.331****	0.019
	(0.061)	(0.060)	(0.062)	(0.068)	(0.064)
AvgDem				−0.034****	−0.048****
				(0.008)	(0.011)
AvgDepend				−0.00001	−0.00002**
				(0.00001)	(0.00001)
BothNuke					0.525
					(0.323)
EitherNuke					0.484***
					(0.148)
Major					1.670****
					(0.152)
CapRatio					−0.225****
					(0.039)
InterestSimilarity					−0.422
					(0.275)
Ally					0.517**
					(0.218)
Time Controls	None	Continuous	Year Dummy	Continuous	Continuous
PeaceYears	Not Included	Included	Included	Included	Included
Geographic Controls	Not Included	Included	Included	Included	Included
Observations	655,241	655,241	655,241	463,370	452,784
Earliest Year	1817	1817	1817	1871	1871
Note:					*p<0.1; **p<0.05; ***p<0.01; ****p<0.001

Table 5.3: Logistic regression *without* dyadic fixed effects. War.

	Dependent variable:				
	War				
	(1)	(2)	(3)	(4)	(5)
AvgLogGDPPC	−0.611****	0.202	0.041	0.294	−0.152
	(0.091)	(0.132)	(0.140)	(0.204)	(0.185)
AvgDem				−0.082****	−0.075****
				(0.020)	(0.021)
AvgDepend				0.00000	−0.00000
				(0.00000)	(0.00002)
BothNuke					−13.661****
					(0.678)
EitherNuke					−0.471
					(0.448)
Major					2.786****
					(0.387)
CapRatio					−0.503****
					(0.079)
InterestSimilarity					−0.860
					(0.813)
Ally					0.100
					(0.628)
Year Controls	None	Continuous	Year Dummy	Continuous	Continuous
PeaceYears	Not Included	Included	Included	Included	Included
Geographic Controls	Not Included	Included	Included	Included	Included
Observations	655,241	655,241	655,241	463,370	452,784
Earliest Year	1817	1817	1817	1871	1871
Note:					*p<0.1; **p<0.05; ***p<0.01; ****p<0.001

Table 5.4: Logistic regression *with* dyadic fixed effects. MIDs.

	Dependent variable:				
	MID				
	(1)	(2)	(3)	(4)	(5)
AvgLogGDPPC	-0.219***	-0.397****	-0.369***	-0.418***	-0.459***
	(0.069)	(0.118)	(0.126)	(0.141)	(0.148)
AvgDem				-0.077****	-0.081****
				(0.012)	(0.011)
AvgDepend				0.000	0.000*
				(0.000)	(0.000)
BothNuke					-0.683**
					(0.270)
EitherNuke					-0.562***
					(0.182)
Major					0.672**
					(0.276)
CapRatio					-0.105
					(0.089)
InterestSimilarity					-0.392
					(0.366)
Ally					0.036
					(0.308)
Time Controls	None	Continuous	Year Dummy	Continuous	Continuous
PeaceYears	Not Included	Included	Included	Included	Included
Geographic Controls	Not Included	Included	Included	Included	Included
Observations	66,154	66,154	66,154	39,623	38,973
Earliest Year	1817	1817	1817	1871	1871
Note:				*p<0.1; **p<0.05; ***p<0.01; ****p<0.001	

Table 5.5: Logistic regression *with* dyadic fixed effects. War.

	Dependent variable:				
	War				
	(1)	(2)	(3)	(4)	(5)
AvgLogGDPPC	-1.151****	-1.678****	-1.455****	-1.055**	-0.687
	(0.114)	(0.356)	(0.431)	(0.495)	(0.466)
AvgDem				-0.187****	-0.188****
				(0.033)	(0.033)
AvgDepend				0.000	0.000
				(0.000)	(0.000)
BothNuke					-13.229****
					(0.317)
EitherNuke					-1.289***
					(0.424)
Major					0.187
					(0.540)
CapRatio					-0.122
					(0.316)
InterestSimilarity					-1.563
					(1.143)
Ally					-0.069
					(0.740)
Time Controls	None	Continuous	Year Dummy	Continuous	Continuous
PeaceYears	Not Included	Included	Included	Included	Included
Geographic Controls	Not Included	Included	Included	Included	Included
Observations	20,097	20,097	20,097	8,128	8,026
Earliest Year	1817	1817	1817	1871	1871
Note:					*p<0.1; **p<0.05; ***p<0.01; ****p<0.001

causes them to fight less and that this unobserved variable is the source of the statistical relationship. It is also possible, that data quality issues or issues with the specification of our model are responsible. However, given how large and robust the relationship is (at least in the case where we focus on MIDs as our dependent variable) I do believe the regression provides substantial evidence for a genuine pacifying effect.

The fact that the relationship does not go away when we include trade and democracy controls also supports the hypothesis that the effect runs through other causal pathways to a substantial extent. The Diminishing Returns Mechanism outlined in Section 5.4, as well as some of the alternative mechanisms outlined in Section 5.5, could play important roles in driving the causal relationship.

5.7 Conclusion

This paper has presented a range of theoretical and empirical reasons to believe that—contrary to the conventional view—economic growth probably lowers the risk of interstate war. The paper makes three main contributions. First, it develops the argument that prosperity supports peace by making people less willing to risk their lives and freedom for financial rewards. Second, it clarifies and synthesizes other possible causal pathways between growth and peace by situating them within a unified bargaining model. Third, using a fixed effects model, it shows that pairs of states become substantially less likely to fight as they grow wealthier.

I have also noted and explored a number of complications and caveats involving the relationship between growth and peace. Most significantly, I found that prosperity can actually raise the risk of war in cases where non-material motives to fight are sufficiently strong. I distinguished between "scarcity-driven wars" (wars that are made more likely by prosperity) and "scarcity-constrained wars" (wars that are made less likely by prosperity) and noted that the latter may be becoming more prevalent over time. This raises the possibility

that growth could eventually begin to have an overall aggravating effect, even if it has mostly had a pacifying effect so far.

Notably, in this paper, I have refrained from making any claims about how the risk of war has changed over time or how it may change in the future. Although this paper certainly *bears* on the question of whether war has declined or will decline, one should be cautious about drawing out any inferences along these lines without careful additional analysis.

5.8 Appendices

5.8.1 Calculating the Approximte Effect of a Doubling in GDPPC

In the logistic regression model used in Section 5.6, the probability of conflict $Pr(\text{Conflict})$ is given by an expression of form:

$$Pr(\text{Conflict}) = \frac{1}{1 + e^{-(X + b \cdot \text{AvgLogGDP})}}$$

Here, b is the coefficient for AvgLogGDPPC and X is a term that is a function of the values of the other variables. We can nearly always assume that the value of $e^{-(X+b\cdot\text{AvgLogGDPPC})}$ is high, since the model nearly always assigns a low probability to conflict. If we make this assumption, then, since $e^{-(X+b\cdot\text{AvgLogGDPPC})} \gg 1$, the following approximation holds:

$$Pr(\text{Conflict}) \approx \frac{1}{e^{-(X+b\cdot\text{AvgLogGDPPC})}} = e^{X+b\cdot\text{AvgLogGDPPC}}$$

Now, suppose that GDP increases by some factor g. This will add $\ln(g)$ to the value of AvgLogGDPPC.

Then the probability of conflict will then change by some factor, R. R can be approximated by:

$$R \approx \frac{e^{X+b\cdot(\text{AvgLogGDPPC}+\ln(g))}}{e^{X+b\cdot\text{AvgLogGDPPC}}} = \exp\left(b \cdot \ln(g)\right) = g^b$$

This means that if GDPPC increases by a factor of g — and the initial risk of conflict is sufficiently low — then we can expect the risk of conflict to change by approximately a factor of g^b.

5.8.2 Rerunning the Regressions with Older GDP Data

The regressions run in this paper use up-to-date estimates of historical GDP-per-capita levels (Bolt and van Zanden, 2020), which run back into the early nineteenth century. Previous studies have largely relied on datasets that contain estimates that are less up-to-date and that cover shorter windows of time. For instance, Gartzke (2007) and follow-up studies largely rely on a dataset that was compiled by Kristian Gleditsch in 2002 and has 1950 as its start year (Gleditsch, 2002).

However, the basic statistical findings in this paper do not seem to depend too much on which GDP-per-capita dataset we use. To show this, I will now rerun the same fixed effect regressions using Gleditsch's widely-used dataset from 2002. (See Tables 5.6 and 5.7.)[39] The fixed effects models continue to suggest a strong association between GDP-per-capita and the absence of militarized interstate disputes, for all sets of controls considered. The relationship between GDP-per-capita and the absence of war *does* become insignificant for all sets of controls (except for only-fixed-effects case); however, since the estimated association is still large, it continues to seem plausible that there is a genuine association that is simply being hidden by large standard errors.

[39]Note that, in the final model in Table 5.7, the control "Major" has been automatically dropped by the feglm function due to collinearity.

Table 5.6: Logistic regression *with* dyadic fixed effects, using GDP-per-capita data from Gleditsch (2002). MIDs.

	Dependent variable:				
	MID				
	(1)	(2)	(3)	(4)	(5)
AvgLogGDPPC2002	-0.311***	-0.440****	-0.408****	-0.573****	-0.559****
	(0.101)	(0.112)	(0.119)	(0.141)	(0.148)
AvgDem				-0.048****	-0.053****
				(0.015)	(0.015)
AvgDepend				0.000	0.000
				(0.000)	(0.000)
BothNuke					-0.949**
					(0.412)
EitherNuke					-0.448**
					(0.190)
Major					1.058**
					(0.484)
CapRatio					-0.330***
					(0.117)
InterestSimilarity					0.111
					(0.475)
Ally					-0.284
					(0.423)
Time Controls	None	Continuous	Year Dummy	Continuous	Continuous
PeaceYears	Not Included	Included	Included	Included	Included
Geographic Controls	Not Included	Included	Included	Included	Included
Observations	32,857	32,857	32,857	24,515	24,515
Earliest Year	1817	1817	1817	1871	1871
Note:					*p<0.1; **p<0.05; ***p<0.01; ****p<0.001

Table 5.7: Logistic regression *with* dyadic fixed effects, using GDP-per-capita data from Gleditsch (2002). War.

	Dependent variable:				
	War				
	(1)	(2)	(3)	(4)	(5)
AvgLogGDPPC2002	-1.715***	-0.832	-1.006	-0.993	-0.908
	(0.388)	(0.610)	(0.677)	(0.613)	(0.645)
AvgDem				-0.176*	-0.161
				(0.089)	(0.092)
AvgDepend				0.000	0.000
				(0.000)	(0.000)
BothNuke					-12.424***
					(0.783)
EitherNuke					-0.655
					(0.637)
CapRatio					-0.974*
					(0.457)
InterestSimilarity					-0.061
					(2.320)
Ally					-0.460
					(1.862)
Time Controls	None	Continuous	Year Dummy	Continuous	Continuous
PeaceYears	Not Included	Included	Included	Included	Included
Geographic Controls	Not Included	Included	Included	Included	Included
Observations	3,467	3,467	3,467	2,456	2,456
Earliest Year	1817	1817	1817	1871	1871
Note:					*p<0.1; **p<0.05; ***p<0.01; ****p<0.001

Chapter 6

Adjusting Our Initial Forecast

6.1 Abstract

This (non-paper) chapter focuses on revising the initial "outside view forecast" that we developed in Chapter 3. To do this, I draw on the understanding of the effect of economic growth that we developed over the course of Chapters 4 and 5. I argue that economic growth could partly explain the downward trend in conflict that seems to have begun in the mid-20th century. Furthermore, if economic growth continues or accelerates, then it could also conceivably produce a large further decrease in the risk of conflict. Ultimately, trying to account for the implications of economic growth does not seem to warrant any radical changes to our initial forecast. This chapter overall provides additional support for the claim that a probability in the low-double digits (e.g. 15%) can be considered a reasonable estimate of the likelihood of a peaceful future.

6.2 Background

6.2.1 Our Outside View Forecast

In Chapter 3, we produced an initial "outside view" forecast of the probability of a peaceful future. This forecast was generated by analyzing and extrapolating historical trends and trying — to the extent possible — to put causal questions to the side.

The outside view forecast ultimately suggested that the probability of a peaceful future is low but not negligible. We gave **20%** as an illustrative example of a probability estimate that could be considered "reasonable," although other estimates that are meaningfully higher or lower could be considered reasonable as well.

The main finding that bolstered the Peaceful Future Hypothesis was that there has seemingly been a fairly steep downward trend in the risk of conflict since around the time of the Second World War. There was substantial evidence of a trend-break around this time. (The best-fit year, under the assumption there were no previous breaks, is 1937.) If this apparent trend continues for along enough, then this would imply a high probability of a peaceful future. However, if we consider the historical frequency of trend breaks, then we should also think there is a substantial chance that the trend will break early and break in the wrong direction. Furthermore, it is still possible that the apparent recent trend is illusory. When we tried to add up all of these considerations — plus some further considerations — we ultimately arrived at the 20% number.

6.2.2 Our Understanding of the Effect of Growth

In Chapter 4 and 5, I explored the effect of economic growth on the risk of conflict. The main upshot is that economic growth probably *does* substantially reduce the risk of interstate conflict.

In this section, I will briefly review the key findings from the relevant chapters. I will also go one step further, however, by generating an explicit probability distribution over the

magnitude of growth's effect.

6.2.2.1 Our Key Statistical Finding

One key finding has been that, as states become wealthier, they become less likely to experience militarized disputes with each other. The relationship is robust to the addition of temporal controls and controls for several other major variables of interest.

Table 6.1 reports the results of the key fixed effect regressions, rerun with "reciprocated use of force" (the variable we focused on in our outside view forecast) as the dependent variable. The regression with the full set of controls suggests that increasing the average GDP-per-capita of a pair of states by some factor R will, all else equal, multiply the risk of conflict by approximately a factor of $R^{-.5}$. For example, the results suggest that: if two states have their incomes double, then the risk of conflict between them will fall by approximately 30%. **The 95% confidence interval for the percentage change in risk is** $[-50\%, -2\%]$ **in the case of a doubling in GDP-per-capita and** $[R^{-.99} - 1, R^{-.03} - 1]$ **in the general case.**[1]

6.2.2.2 What Does This Statistical Finding Imply?

Importantly, the confidence interval given above should not be confused with a "credible interval." The numbers contained within a "confidence interval" are numbers that — if they were the parameter's true value — would *have a high chance of generating the data-points we have observed.* On the other hand, the numbers contained in a "credible interval" are numbers that *have a high chance of being the parameter's true value.*

As a toy example, to illustrate this point, suppose that we are initially *extremely confident* — let's say, 99.999% confident — that some coin is evenly weighted and thus has

[1]When the regression is run with a more limited sets of controls, however, the estimated effect is substantially larger. For instance, if we include only temporal and geographic controls, we find that the point estimate and confidence interval for the effect produced by a doubling become -43% and $[-58\%, -24\%]$ respectively. This is worth at least briefly noting, since — as discussed in the previous chapter — we should not necessarily assume that the estimate produced using the full set of controls is the most accurate one.

Table 6.1: Logistic regression *with* dyadic fixed effects. Recripocated use of force.

	Dependent variable:				
	Reciprocated Force				
	(1)	(2)	(3)	(4)	(5)
AvgLogGDPPC	-0.650****	-0.850****	-0.818****	-0.573**	-0.506**
	(0.094)	(0.202)	(0.214)	(0.258)	(0.244)
AvgDem				−0.035****	−0.048****
				(0.008)	(0.011)
AvgDem				-0.079****	-0.089****
				(0.021)	(0.022)
AvgDepend				0.000***	0.000**
				(0.000)	(0.000)
BothNuke					0.079
					(0.548)
EitherNuke					-1.274****
					(0.278)
Major					0.196
					(0.425)
CapRatio					-0.279*
					(0.162)
InterestSimilarity					0.019
					(0.489)
Ally					-0.294
					(0.398)
Time Controls	None	Continuous	Year Dummy	Continuous	Continuous
PeaceYears	Not Included	Included	Included	Included	Included
Geographic Controls	Not Included	Included	Included	Included	Included
Observations	36,309	36,309	36,309	18,419	18,089
Earliest Year	1817	1817	1817	1871	1871
Note:					*p<0.1; **p<0.05; ***p<0.01; ****p<0.001

a 50% chance of landing heads each time it is flipped. (We made the coin ourselves, we have examined it physically, we have made thousands of other coins in the same way and they have all turned out to be evenly weighted, and so on.) We flip the coin 100 times and it lands heads 60 times. The 95% *confidence interval* would be [.504,.696] and just barely exclude the weighting of .5, since an evenly weighted coin is unlikely to land heads so many times. However, our 95% *credible interval* would simply be [.5,.5], since — although we should now be *less* confident that the coin is evenly weighted — we should nonetheless still think there is at least a 95% chance that it is evenly weighted. Since we began with such a high level of confidence that the coin is evenly weighted, our leading hypothesis should be that we simply observed an unlikely sequence of flips.

The point is: Credible intervals are conceptually distinct from confidence intervals and can diverge from them.

6.2.2.3 Producing a Credible Interval

To forecast how economic growth will change the rate of conflict, we ultimately need to produce a credible interval over the magnitude of its effect. In some cases, it is reasonable to directly base the values of one's credible interval on the confidence interval produced by a regression.[2] However, we do not seem to be in one of those cases. Our credible interval at least should be different from the confidence interval.

One core issue is that we are uncertain about how well the statistical model is actually controlling for all relevant confounders. Although we may have controlled for a substantial portion of omitted variable bias, through our use of two-way-fixed effects, the model still certainly does not entirely control for everything that matters and thus is almost certainly producing a biased estimate. Unfortunately, we do not know in what direction the bias runs. This consideration suggests that our credible interval should be wider interval than

[2]Specifically, it is reasonable to do this if (a) we expect the regression to produce an unbiased estimate of the parameter's true value and (b) we have no further evidence concerning the parameter's true value and no prior reason to consider any value more-or-less likely than any other.

the confidence interval.

We also, of course, have access to additional arguments and evidence that should influence the probabilities we assign to different effect sizes. Chapter 5 provided a fairly wide-ranging tour of some of these arguments and evidence.

Overall, key considerations that should influence our estimate include:

- We know that effect sizes implied by regressions on observational data are often highly misleading, perhaps especially in the social sciences, due to bias from confounders that are not fully controlled for.

- Our regression included dyadic fixed effects controls, as well as temporal controls, which should quite substantially reduce bias. However, these controls certainly do not suffice to *completely* remove bias.

- Our full regression includes controls for some additional control variables — particularly, economic interdependence and democracy — that may actually *in part* mediate the effect of growth on conflict. Since including these controls meaningfully reduces our effect size estimate, it is possible that the estimate we have produced using these controls is biased at least slightly downward. More generally, we should not assume that the regression with the full set of controls is in fact producing the least biased estimate.

- The data we are using in our regressions – for our independent variable as well as for our control variables — certainly is not entirely reliable. The one-dimensional metrics we are using to quantify certain inherently multi-dimensional traits, for instance the Polity2 scores we are using to measure how democratic a state is, are also all unavoidably crude and arbitrary to some degree. We should not believe, for example, that we are genuinely fully controlling for variation in democracy or capability gaps by including the control variables "AvgDem" and "CapRatio."

- We have found that there are *several* different theoretically plausible mechanisms by which growth could substantially decrease the risk of conflict. There is also empirical evidence for the validity of some of these mechanisms, such as the "Diminishing Returns Mechanism." These findings lend additional plausibility to the hypothesis that growth has a substantial pacifying effect.

- However, some of the mechanisms — including the Diminishing Returns Mechanism — suggest that the effect of growth may decrease as states become wealthier. This means that the effect may now be lower than it has been throughout the past.

- There are also some mechanisms that could lead growth to increase the risk of war, at least under certain circumstances. For example, we found that for some "scarcity-constrained wars" prosperity can exacerbate risk by make people more willing to tolerate financial losses for the sake of achieving immaterial (e.g. ideological or ethical) objectives.

- More generally, considering a range of possible causal mechanisms makes it clear that the causal link between growth and conflict is likely to be extremely complex, difficult to entangle from other causal relationships, and difficult to reason about. Analysis of most of the relevant mechanisms has also been extremely cursory. Our theoretical understanding of the relationship between growth and peace is therefore probably still fairly poor — and we should therefore be wary of placing too much stock in theoretical arguments.

- We also know that, in the social sciences, the effects of changes in individual variables are often small. Nonetheless, since growth is such a transformative phenomenon, it seems *a priori* particularly plausible that its effect will be unusually large.

Together, in my view, this set of considerations suggests a credible interval that: (a) is wider than the confidence interval produced by the regression, (b) assigns a non-negligible

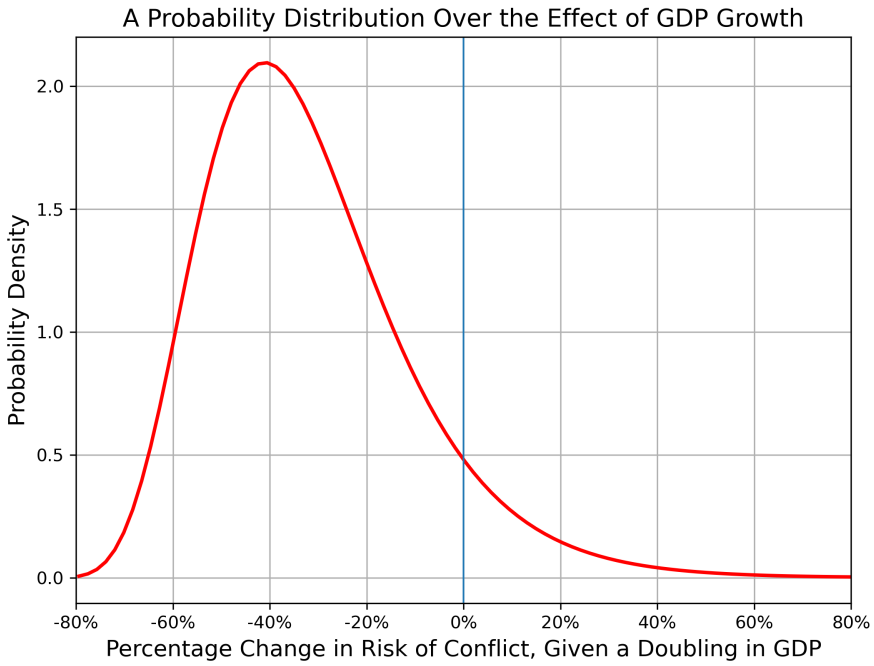

Figure 6.1: A probability distribution over the effect size of growth. Here "conflict" is operationalized as the outbreak of a dispute involving the reciprocated use of force.

probability to growth actually having an aggravating effect, and (c) still implies that it is overall more-likely-than-not that growth has a large pacifying effect.

In principle, we could develop a formal Bayesian model to properly aggregate all of the relevant considerations and produce an overall credible interval for the effect of growth. However, I will again take a crude-but-quick approach and simply report a credible interval that satisfies the three conditions listed above.

Specifically: Let us suppose, again, that the GDP of some pair of states doubles. *I then will say that one reasonable credible interval, for the proportional effect that this has on the risk of war, is [-65%, 20%].*

Figure 6.1 plots the implied probability over the effect of growth, assuming a log-normal distribution.

As noted above, however, it is plausible that the effect of growth will decline further as the world becomes wealthier. The Diminishing Returns mechanism, in particular, implies

that this is likely. More generally, across a range of domains, it is common for the marginal effects of variables to fall off once they reach sufficiently extreme values. This means that it would be natural for us to adjust this probability distribution if we consider larger and larger numbers of doublings. If we were to consider two doublings, for example, then it would seemingly make sense to at least slightly compress the probability distribution and shift it at least slightly to the right.

6.3 Growth and the History of Conflict

Now that we have laid out our outside view forecast, and laid out our estimate of the effect of growth, we can move toward synthesizing them. A natural first step is to consider whether growth could, at least in part, explain the historical trends that have informed our outside view forecast.

6.3.1 Can Growth Explain the (Apparent) Downward Trend in Conflict?

A downward trend in the rate of conflict seems to have started around the time of the Second World War.

Could economic growth help to explain the emergence of this trend? I think this explanation — although surely incomplete — is quite plausible. For reference, consider Figure 6.1, which again plots the GDP-per-capita of the median state included in the Correlates of War dataset over the past two centuries. The rate of growth was clearly far higher in the post-war period than the pre-war period. The median state only really began to experience meaningful growth in the 1950s.[3]

[3]Of course, the timelines here do not perfectly line up. If we assume there were no previous trend breaks, then the best-fit year for the trend break is 1937. If we assume there was at least one previous trend break, which seems about equally likely, then the best fit year is instead 1944. These dates are both somewhat earlier than the point where the median states to properly grow.

However, it is not clear that we get hung up on specific best-fit years. Since nearby years may provide

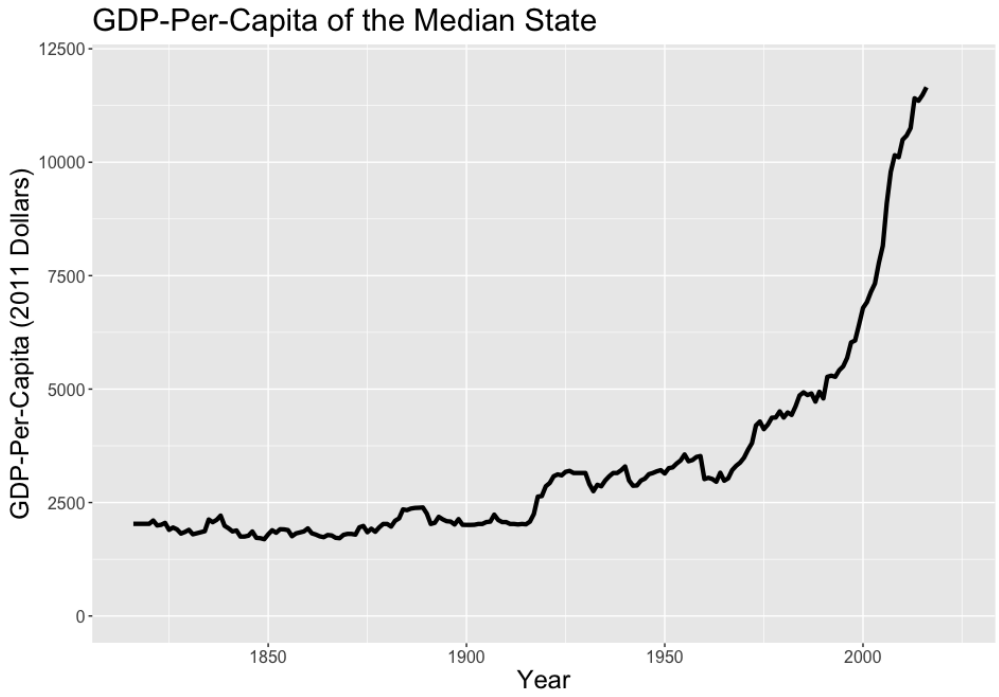

Figure 6.2: The GDP-per-capita of the median state included in the Correlates of War dataset, according to estimates in the 2020 Maddison Project Database (Bolt and van Zanden, 2020). Note, as before, I have interpolated missing values for individual years for individual countries.

The estimated effect size of growth is also large enough to account for most of the decline. The best-fit parameters for our statistical model imply that the expected rate of conflict fell by about a *factor of six* between 1950 and 2010. Meanwhile, the data plotted above shows that the median state became approximately four times wealthier over the same period. To explain this decline in the expected rate of war, then, each doubling in GDP-per-capita would need to decrease the risk of conflict by approximately 60%.[4] This effect size falls substantially below our median estimate of approximately -40% (see Figure 6.1), but is just barely inside our credible interval of [-65%,+20%].[5]

similarly good fits, we should not assume that trend break occurred at precisely these dates. Indeed, if we run a regression on the period from 1937 to 1950, then — unsurprisingly — we do not find statistically significant evidence that the trend existed during this period. The notion of a discrete break that is localized to a single year is, in any case, presumably something of an idealization.

[4]In general, if GDP increases by a factor of g, and the risk of war changes by a factor of R, then the implies that the risk of war is multiplied by a factor of $R^{ln(2)/ln(g)}$ at each doubling.

[5]Of course, especially since we needed to make debatable judgement calls to arrive at this particular credible interval, one should not make too much of the fact that a number is just inside or just outside of the interval.

In any case, since GDP-per-capita is far from the only variable that influences the rate of conflict, we do not need to assume that is the only variable that could be contributing to a downward trend in conflict. Our median estimate of growth's effect is still large enough to account for about *four-fifths* of the decline. Then, as we approach the upper end of our credible interval, the effect actually becomes large enough to account for *all* of the decline.

Overall, the apparent downward trend in conflict could be explained in large part by the pacifying effect of economic growth.

6.3.2 Other Interpretations

Of course, there are other possible interpretations of historical trends. I will now briefly walk through three alternative possibilities: (1) the possibility that the downward trend has mostly been caused by a shift in international order, (2) the possibility that the downward trend has mostly been caused some other factors, and (3) the possibility that there actually is no downward trend in the rate of conflict.

6.3.2.1 The Downward Trend Has Mostly Been Caused by a Shift in International Order

At the end of the Cold War, world politics shifted from a bipolar structure to a unipolar structure. There was now only a single superpower, whereas previously there had been two.

A substantial literature in International Relations suggests that unipolar world order may be more conducive to peace than bipolar world order (e.g. Wohlforth, 2014). In a bipolar world order, a substantial portion of conflicts may arise out of the two superpowers' competing efforts to exert global influence. In a unipolar world, on the other hand, the single superpower may be able to serve as a reliable mediator or enforcer of peace across much of the world.

Braumoeller (2019) argues that this shift in polarity explains why the rate of conflict

is now lower than it was at the start of the Cold War. This is a reasonable hypothesis. It overall strikes me as quite plausible that the shift has at least been a contributing factor in the recent decline of war.

However, **the post–Cold War shift in polarity does seem like it can be the *primary* explanation for the downward trend we have observed.** The clearest way to see this is to run separate regressions just on the years of the Cold War and just on the years after the Cold War. As Table 6.2 reports, we find statistically significant evidence of a downward trend in the rate of conflict both during and after the Cold War. This means that there was not simply a single jump to a lower rate of war at the end of the Cold War. Rather, there has seemingly been a downward trend across the full post-WWII period.

Notably, the trend does seem to have been much steeper after the end of the Cold War. One natural explanation for this finding is that the median state's growth rate was actually much higher in the post–Cold War period than it was in the pre–Cold War period. (The average growth rate was approximately 1% between 1945 and 1989, but approximately 3% between 1989 and 2014.) It is also possible, though, that the end of the Cold War bolstered some additional pacifying trends in important ways.

Nonetheless, it does not seem to be the case that the end of the Cold War is the only reason or even the *main* reason why the rate of conflict has been declining.

6.3.2.2 The Downward Trend Has Mostly Been Caused by Something Else

Of course, since the mid-20th century, the world has changed in a large number of ways. Although changes in polarity and prosperity are surely two of the most salient changes, it would also be worth exploring whether (e.g.) changes in the level of economic interdependence or the prevalence of democracy could independently play major explanatory roles.[6] It is also conceivable that less-material trends, such as changes in values or norms, could play important explanatory roles as well.

[6]As has been discussed, growth likely supports both interdependence and democracy to some extent. However, interdependence and democracy also can and do increase on their own.

Table 6.2: Time trends before and after the end of Cold War.

	Dependent variable:	
	Log(CurrentRate/PreviousRate)	
	(1)	(2)
	$1945 \leq Year \leq 1989$	$1989 \leq Year$
YearsSince1816	-0.016^{**}	-0.063^{***}
	(0.007)	(0.018)
Log(PreviousRate)	-0.783^{***}	-1.083^{***}
	(0.117)	(0.212)
Constant	-0.808	6.480^{**}
	(0.867)	(2.916)
Observations	45	22
Note:	*p<0.1; **p<0.05; ***p<0.01	

However, I must leave these explorations for later work. The project of comprehensively studying all plausible pacifying trends — in the same depth with which we have studied growth — would be a time-intensive one, although I also think a worthwhile one.

6.3.2.3 There is No Downward Trend

As discussed in Chapter 3, our finding that there is has been a downward trend in the post-1937 period is highly statistically significant ($p < 10^{-6}$). Since the data we have for the post-war period is also probably fairly reliable and comprehensive, this finding is also unlikely to be an artifact of data quality issues.

The finding does depend, though, on our choice of model. If we had fitted some substantially different model to the data, then it is possible we would have failed to find a significant result. Although the broad class of models we have focused on strikes me as the most obvious one — and I don't see any obvious alternative that are both similarly simple and would fit the data similarly well — it is hard to rule out the possibility that some other

class of models would be superior. There are also various small modifications we could make to our models, such as allowing for error terms that are not normally distributed, that would certainly lead to at least somewhat different significance estimates.

The statistically significant result was also produced, it is important to note, by a variant of our model that allows for trend breaks. Since this variant gives us the freedom to search for evidence of a downward trend after each possible breakpoint, we may be able to find a highly significant result even if there is not actually any underlying trend.[7]

In Chapter 3, we used the Bayesian Information Criterion (BIC) as a way of weighing up different variants of our core model that we could fit to the data. Among variants that allow for time trends, the criterion suggested that variants with trend breaks are indeed substantially more likely to be approximately "correct" than variants without trend breaks. However, the criterion also resulted in us assigning about a 50% probability to the simple "null model" — which does not allow for time trends at all — being correct. This would suggest there is only about a 50% chance that there is actually a downward trend.

However, these BIC-derived relative likelihoods are based on the assumption that we have no prior reason to think that one model is more likely than another. In fact, the analysis we have just conducted — which suggested that growth is likely to have had a large pacifying effect over the past several decades — does give us a fairly compelling reason to downweight the null model. If we *start out* giving a quarter as much weight to the null model, for example, then the BIC would now imply that the null model has only a 13% probability of being correct.[8]

Overall, I do think is quite likely that there has been a genuine post-1937 trend in the rate of conflict. However, we still cannot entirely rule out the possibility that the trend

[7] This concern here is *overfitting*.

[8] Another point to keep in mind — which was also noted in Chapter 3 — is that the use of the BIC to produce relative probabilities of "model correctness" is somewhat controversial (Gelman et al., 2014). Some critics charge that it excessively privileges simple models. Other critics charge that it is based on questionable theoretical assumptions or, more fundamentally, that it is simply misguided to talk about "the probability that a model is correct" given that no model is entirely correct. Nonetheless, when it comes to the task of assigning likelihoods to trend breaks, the BIC still tends to be the mostly commonly used tool.

is illusory.

6.3.3 Summary

Economic growth is a compelling explanation for the downward trend in conflict that seems to have begun in the mid-20th century. The median state began to experience substantial economic growth at around the time the trend began. The estimated effect of growth is also large enough to account for the majority of the decline in the expected rate of conflict.

We should keep in mind that there is still some chance that the downward trend is illusory. However, overall, it does seem to be real. The end of the Cold War and the corresponding shift in polarity cannot account for the trend, since the trend seemingly existed both before and after the end of the Cold War. It is possible that *other* causal factors would offer more compelling explanations for the trend. But I will leave a thorough study of alternative explanations for other work.

6.4 Growth and the Future of Conflict

Now, let's turn to considering the implications of economic growth for the future of conflict. The aim will be to understand whether economic growth could plausibly a large role in producing a peaceful future. If we find that it can, then this would reinforce our initial trend-based view that a peaceful future is at least plausible.

We can structure our analysis by breaking it down into three stages:

1) First, we judge how likely different levels of economic growth are in the coming decades.

2) Second, for each potential level of economic growth, we judge how likely it is that this level of growth would be sufficient to cause a dramatic (i.e. $> 90\%$) decline in

the expected rate of conflict.

3) Finally, we aggregate these probabilities to calculate the overall probability that economic growth in the coming decades will be sufficient to cause a dramatic decline in the expected rate of conflict.

The final probability produced by this exercise is a kind of *conditional forecast*. It roughly corresponds to the likelihood that there would be a peaceful future, conditional on the assumption that the net effect of changes in other causal factors will be negligible.[9] For example, we are assuming that the effect of economic growth will not be washed away by the much larger aggravating effect of some other change in the world.

I will discuss both the value and the limitations of this kind of conditional forecast below. Ultimately, we will use it to reassess and update the initial "outside view" forecast that we produced in Chapter 3.

6.4.1 Forecasting Growth

As we in the case of conflict forecasting, we can apply both "outside view" approaches (i.e. approaches based on the analysis of historical frequencies and trends) and "inside view" approaches (i.e. approaches based on causal analysis) when forecasting economic growth over the next century.

To take an outside view approach, we can attempt to forecast economic growth using time series models to extrapolate growth trends forward. There is a large literature on the use of time-series forecasting methods for economic forecasting.

To take an inside view approach, on the other hand, we can use "structural models" to forecast growth. These models draw on economic theories to describe economic output as a function of underlying economic variables that evolve over time. There is also a large

[9] In fact, it does not perfectly correspond to this probability. Since the rate of conflict evolves in such a random and erratic way, the observed rate of conflict within a period can be very different from the expected rate. However, my expectation is that this forecasting method should produce at least *a very rough approximation* of the likelihood of a conditional future.

literature on the use of structural models for economic forecasting. In addition to drawing on formal models, we can also consider more informal arguments or considerations about factors that could boost or stand in the way of growth.

Since we are attempting to produce only a fairly crude forecast here, and since I am — frankly — not in fact an expert in economic forecasting, I will not delve into the complexities of different economic forecasting approaches. Instead, I will provide only a quite rough analysis of the options.

6.4.1.1 Choosing a Variable to Forecast

In principle, if we want to be fully rigorous, we should actually attempt to forecast growth within all individual states and account for correlations between their growth rates. However, to avoid all of the difficulties and complexities involved in doing this, I will make the simplifying assumption that all states will grow at roughly the same rate. This allows me to focus on forecasting growth in a single variable: global GDP-per-capita. (See the appendix of this chapter for a brief comment on the ways in which this simplifying assumption may be distorting.)

6.4.1.2 Choosing Scenarios

As another useful simplification, I will break down possible levels of growth into three broad categories. I will also focus on the next hundred years, which brings us mid-way into the period (2050-2300) covered by our formulation of the Peaceful Future Hypothesis.

- **Steady Growth:** For at least a substantial portion of the next two centuries, growth does not change very much from its current rate. The result is that economic output ultimately becomes multiple times (e.g. 2-20 times) higher than its current level.
- **Low Growth:** Growth substantially declines or even reverses in the coming decades. The result is that economic output at increases at most modestly (e.g. < 2 times) from its current level.

- **High Growth:** At some point, for an extended stretch of time, the rate of growth increases substantially. The result is that economic output becomes vastly (e.g. 20+) higher than its current level.

The steady growth scenario would imply that growth could conceivably bring peace, if its marginal effect is and continues to be *large*. The high growth scenario would imply that could growth could conceivably bring peace even if marginal effect is and continues to be merely *non-negligible*. The low growth scenario, on the other hand, would imply that we have no reason to expect growth to produce substantially less conflict.

6.4.1.3 Outside View Analysis

Figure 6.3 presents a log-scale graph of estimated world GDP-per-capita over the past 200 years. Growth has seemingly been *relatively* steady over the two centuries. Although it was meaningfully slower before the middle of the twentieth century, it has been especially stable since then. Any simple approach to time series analysis, based on this data, will imply a high likelihood of "Steady Growth."[10]

However, there are also other datasets we can use that would suggest somewhat different predictions. In particular, very long-run datasets (extending back before 1820) tend to suggest a trend break around the start of the nineteenth century. Economists generally accept that there was no substantial, sustained upward trend in global per-capita wealth until the Industrial Revolution. The existence of a trend break two hundred years ago — which corresponded to a substantial increase in the growth rate — should lead us to assign a non-negligible probability to another major trend break in the coming centuries. If we give some weight to possibility of a major trend break, then, it seems, we should also give some weight to the "Stagnation or Contraction" or "Acceleration" possibilities.

Some analysts who have studied very long-run datasets of growth estimates (e.g. going

[10]Notably, as shown in Figure 5.1, the median state's growth has been far less steady. This discrepancy follows from the fact that the phenomenon of per-capita growth a small number of states and only gradually spread across the world. The median state actually had fairly negligible growth before the mid-20th century.

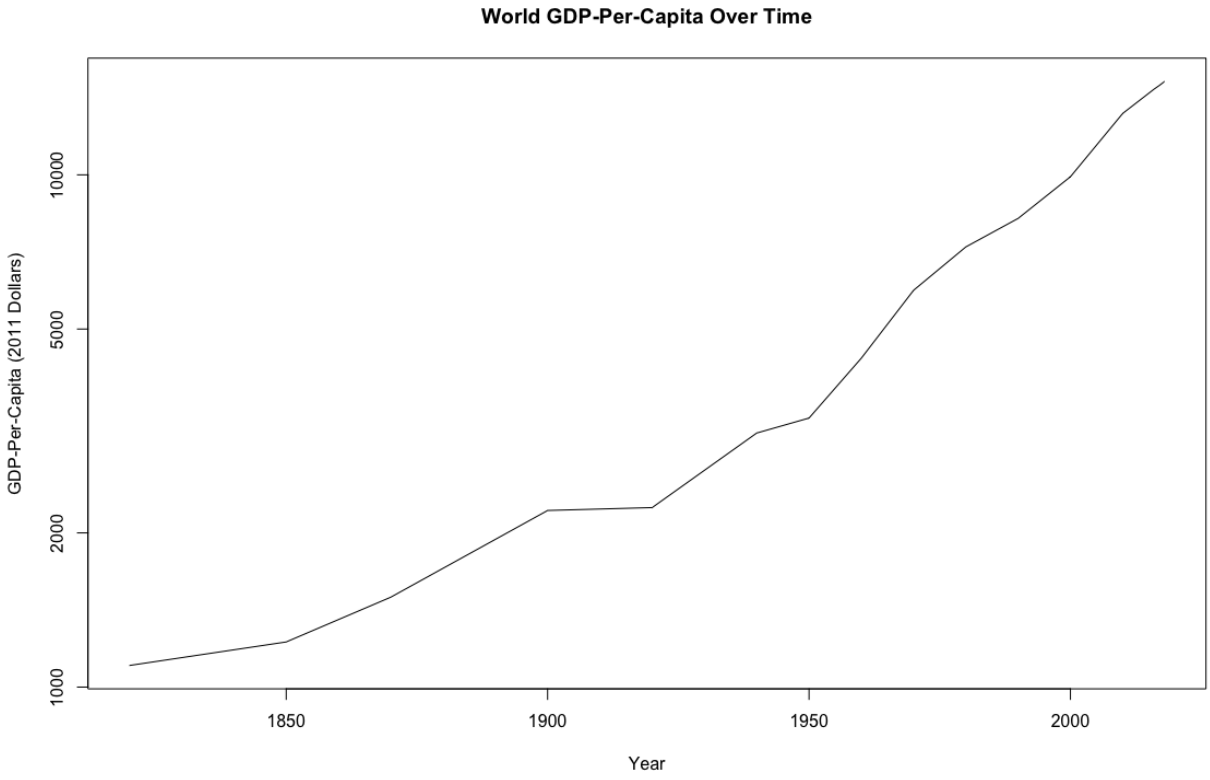

Figure 6.3: Log-scale plot of world GDP-per-capita over the past two centuries, according to data from the Maddison Project (Bolt and van Zanden, 2020). Steady exponential growth would appear as a straight line. We can see that the average global income has been (somewhat) steadily increasing for about two hundred years. (However, as noted above, the gradual spread of growth across the world led the median state to only begin growing substantially in the mid-20th century.)

back more than 500 years) have also engaged in trend extrapolation exercises that suggest a high probability of a major growth increase in the next century. For example, Roodman (2020) models that modern growth increase not as a trend break, but rather the continuation of a long-term — but noisy — upward trend in growth rates. His statistical model assigns a very high likelihood to a major growth increase within the next century. Similarly, Hanson (2000) takes a very long-run perspective and models growth in total output over the past ten thousand years as passing through a series of exponential modes with increasing growth rates. (The growth rate for total output also seems to have increased substantially over the course of the transition to agriculture.) Based on this pattern, he assigns a high likelihood

to the hypothesis that the world will transition to another even faster growth mode in the coming centuries.

However, the same long-run data also admits very different interpretations. Gordon (2017), predicts stagnation in part based on the observation that the past century-and-a-half of per-capita growth has been far greater than the historical average. Insofar as economic stagnation is typical, he argues, we should not be too surprised if stagnation resumes within the next century. While he does not formalize his argument through a statistical model, some simple models — for instance, a model that assumes that growth rates fluctuate erratically around an unchanging baseline — would imply that growth rates are likely to move back toward their (low) historical mean.

Overall, it seems reasonable to say, we should give most weight to the "Steady Growth" scenario. This is what trend extrapolation methods applied to the past two hundred years of data will overwhelmingly predict. However, the long-run picture — along with, in general, the point that we can never fully rule out trend breaks — should also lead us to assign some weight to the "High Growth" and "Low Growth" scenarios.

6.4.1.4 Inside View Analysis

6.4.1.4.1 Arguments for Steady Growth Typical structural models designed for growth forecasting will tend to point to relatively steady (or perhaps moderately declining) growth. Commonly used structural models — for instance, by the International Monetary Fund — account for a range of casual factors such as interest rates and public debt. However, there is no obvious reason to expect these kinds of factors to change in ways that dramatically alter growth rates over the long term. Indeed — although though it is rare to use these kinds of models to forecast growth rates more than a decade out — the few long-term model-derived growth estimates that do exist tend to suggest that global growth rates will not change very much. For instance, PwC has produced a set of national growth estimates up until 2050 that — although they certainly involve variation in growth rates — do overall

correspond to the "Steady Growth" scenario (PwC Economics, 2013).

The most basic class of growth models used by growth theorists ("exogenous growth models") also tend to imply steady-state growth. Specifically, these models imply that growth rates will remain steady into the indefinite future so long as key parameters (e.g. investment rates) do not meaningfully change (Jones and Vollrath, 2013).

There are also, more informally, various causal considerations that can be brought to bear to argue against the possibility of a meaningful increase in the growth rate (Davidson, 2021). For example, some scholars have argued that — due to energetic constraints or other bottlenecks to growth — the world may not be far from the maximum practical rate of economic growth (Smil, 2019).

6.4.1.4.2 Arguments for Low Growth One significant warning sign for the future of growth — noted by growth skeptics such as Gordon (2017) — is that productivity growth rates in most developed countries have been declining for decades. Basic economic theory predicts that, in the long run, it is impossible to sustain per-capita output growth without growth in productivity (i.e. growth in the efficiency with which economic inputs can be transformed into outputs).

There is also evidence that productivity-enhancing innovation is becoming increasingly difficult, in many domains, as the low-hanging fruit are picked away (Bloom et al., 2020). This "fishing out effect" can be partly offset by rising levels of R&D effort, for some time, but the number of researchers and engineers cannot continue to grow at the current rate indefinitely; the current rate is substantially higher than the population growth rate, and, furthermore, the population growth rate is currently declining. A fall in the rate of growth of R&D is one possible explanation for why the economic growth rate could fall in the coming century (Jones, 2023).

Furthermore, there are several factors that could even conceivably lead to large-scale economic backsliding (MacAskill, 2022). A sufficiently extreme global catastrophe, such

as nuclear war or worst-case climate change, could set the world back. The depletion of certain finite natural resources, such as coal, could gradually cripple the global economy if substitutes are not developed and adopted quickly enough. Harmful political trends could also lead governments to engage in more consistent and severe economic mismanagement. Although none of these mechanisms seems likely, in an absolute sense, they do provide us with some reasons to entertain the possibility of backsliding.

6.4.1.4.3 Arguments for High Growth Recently, a number of leading growth economists — including the Nobel Laureate William Nordhaus — have begun to explore whether a growth acceleration could be produced by progress in artificial intelligence (Aghion et al., 2018; Nordhaus, 2021, Trammell and Korinek, 2023). Most AI researchers believe that AI systems are likely to be able to match all economically relevant human capabilities within the next century (Zhang et al., 2022). Exploring the implications of this possibility, these growth economists have noted that some standard models imply that growth rates will accelerate wildly if human labor starts to become largely unnecessary in economic production.

Expressed in simplified terms, the core idea is that building and inventing machines currently requires both "labor" (i.e. human workers) and "capital" (e.g. other machines). Unfortunately, the supply of human workers changes only fairly slowly and is limited by birthrates. The supply of human workers is therefore a bottleneck on production and invention. If human workers largely stopped being needed, though, then this particular bottleneck would be removed. Machines could build more machines and invent better machines, which could then build even more machines and invent even better machines, and so on, in a comparatively unconstrained feedback loop. Economic output, the argument goes, could increase at a higher rate than it does in the current economic regime.

If we suppose, for instance, that the average growth rate over the coming century could be 8% — due to a major surge at some point during the century — then this would be enough to make output-per-person at the end of the century more than 2000 times greater. The

average personâs income would then be equivalent to that of a present-day multimillionaire. Even a much more modest increase in the average growth rate, to 4%, would be enough to increase output-per-person by approximately a factor of 50.

Overall, the validity of the relevant growth models is unclear. It is easy to imagine, for example, that practical constraints that the models do not include (e.g. constraints related to energy availability or to regulation) would mitigate the accelerant effect of automation. It is also unclear just how likely the necessary AI progress is to occur in the coming century.

However, this line of argument does provide us with an additional reason to give the "High Growth" hypothesis some weight. It is also possible to imagine that other technological changes — for instance, breakthroughs in energy technology in the coming century — could offer alternative paths to higher growth rates.

6.4.1.5 Assigning Likelihoods

Clearly, generating an overall growth forecast requires exercising a significant degree of individual judgement. Different people could look at the relevant trends and the relevant arguments and ultimately arrive at fairly different forecasts.

However, overall, it still seems reasonable to me to give the majority of weight to the steady growth scenario. This is what the most obvious approaches to trend extrapolation would suggest. The arguments for the low growth and high growth scenarios also do not seem strong enough, relative to the arguments for steady growth, to override this prediction.

One 2014-2015 survey of economists also suggests that steady growth is the scenario that, on average, experts give the most weight to (Christensen, 2018). The combined expert forecast implied a median estimate of a *six-fold* increase in GDP-per-capita by between 2010 and 2100, which corresponds to steady growth. At the same time, though, the 10th and 90th percentile estimates correspond to an approximately *two-fold* increase (the upper bound for "Low Growth") and an approximately *twenty-fold* increase (the lower bound for "High Growth").

Overall, it does also strike me as defensible to assign *slightly* higher probabilities for the alternative scenarios than the survey suggests. One for assigning slightly higher probabilities is that I expect that this sample of expert would not have produced their estimates while trying to incorporate more recent arguments about implications of AI progress or possibilities related to global catastrophes; in general, due to the myriad forms of uncertainty we face, it strikes me as reasonable to have very wide bounds on what could happen over the next century. For example, assigning something like a **60% probability to steady growth**, a **20% probability to low growth**, and a **20% probability to high growth** also strikes me as still within the "realm of defensibility."

6.4.2 Implications of Different Growth Scenarios

We have just assessed the *likelihoods* of different growth scenarios. The next step in this forecasting exercise is to assess the *implications* of the different scenarios.

To carry out this next step, we can move through each scenario in turn and consider its implications. For each, we can consider how large the pacifying effect of growth would need to be to cause the rate of war to decline dramatically.

We should make sure to account for the possibility that growth has already been contributing to the decline of war. As a crude way of doing this, we can use GDP-per-capita in 1950 (approximately $3,000 in 2011 dollars) as the baseline that we compare future GDP-per-capita levels against. Relative to this baseline, GDP-per-capita has already increased by about a factor of *four*. This meaningfully reduces the amount of additional growth that may be needed to account for a large change in the risk of war.

For the low growth scenario, the effect of growth would clearly need to be exceptionally large to cause the risk of war to decline by anything close to a factor of ten. We can therefore assign an approximately **0% probability** to the possibility that low growth would be sufficient to produce a major growth-caused decline in war.

For the steady growth scenario, the effect of growth would need to be very large but

not quite as exceptionally large. To pick out a particular level of growth, to represent this scenario, let us suppose that there will be a 5x increase in per-capita output from its current level. This implies a 20x increase from its level in 1950. For the expected rate of war to decline by 90%, then, each doubling in output would need to reduce the rate of war by roughly 60%. The credible interval we chose above — in Section 6.2.2.3 — then suggests that there may be an approximately **6% chance** that the effect would be sufficiently large.

For the high growth scenario, the effect of growth could be somewhat smaller. To pick out a particular level of growth, to represent this scenario, let us suppose that there will be a 100x increase in per-capita output from its current level. This becomes 400x once we take into account the decline since 1950. For the expected rate of war to decline by 90%, each doubling in output would, on average, need to reduce the rate of war by roughly 24%. The same credible interval suggests there is something like a 70% chance the effect would be sufficiently large. However, to account for the possibility that the marginal effect of growth will decline as the world becomes vastly wealthier, it seemingly makes sense to use a lower figure. Fairly arbitrarily, then, I will cut the probability estimate in half and bring it down to a **35%** chance that the effect will be sufficiently large.

6.4.3 A Conditional Forecast

Table 6.3 reports the values of the key probabilities we have estimated. Once again, these values should be considered *illustrative* to a certain degree, since reasonable people could report meaningfully different values. However, now that we have the values, we can combine them to reason about the overall likelihood that economic growth will cause the rate of conflict to substantially decline. As Table 6.3 reports, the relevant variables multiply through to imply an **11% chance** that the risk of war will decline by at least a factor of 10.

Since we have been focusing only the effect of growth — which is just a single causal factor, out of many, that could plausibly influence the rate of conflict — this forecast is a "conditional forecast." We are forecasting the chance of a decline in conflict *conditional*

Table 6.3: Calculation of the probability that economic growth will be sufficient to produce an overall 90% decline in the expected rate of conflict (assuming that other changes in the world will have no net effect).

Growth Scenario	**Probability of Scenario**	**Probability of Sufficient Effect**	**Contribution to Probability of Sufficient Risk Reduction (Product of Columns)**
Low Growth	20%	0%	0%
Steady Growth	60%	6%	4%
High Growth	20%	35%	7%
Total Probability of Sufficient Risk Reduction (Sum of Rows)			11%

on the assumption that net effect of changes in other causal factors will be negligible. This assumption is almost certainly wrong.

Nonetheless, the conditional forecast is still informative. It is true that causal factors that could have aggravating effects that offset the pacifying effect of growth. On the other hand, other causal factors could also have pacifying effects that add to the effect of growth. If we do not have a strong reason to expect other factors to be aggravating, on net, then the conditional forecast is sufficient to suggest that there is a non-negligible chance of a peaceful future. We would need to be very confident that the other factors will be very aggravating, for example, to justify cutting the probability estimate in half.

6.5 Conclusion: A New Overall Forecast

To recap, in Chapter 3, we produced an "outside view" forecast that suggested that a peaceful future is unlikely but not implausible. Specifically, as semi-illustrative value, we assigned a 20% probability to a peaceful future. The main finding that bolstered this probability was that a substantial downward trend in the rate of conflict seems to have begun in the mid-20th century.

We then considered whether economic growth could, at least in part, explain this apparent downward trend. We found that it could. As explored in Chapters 4 and 5, economic growth seems to have a large pacifying effect. The typical state's growth rate also became much higher at about the right time. The fact that we have identified a plausible causal explanation for the apparent downward trend should increase our confidence that the trend is real.

We next considered whether growth could plausibly lead to a large enough decline, over the next century, to produce a peaceful future. We found that — if we assume that other causal factors will have a negligible net impact on the risk of conflict — then growth alone could plausibly produce a peaceful future. Specifically, although the specific number here is somewhat arbitrary, we found that growth alone would seemingly have about a 10% chance of producing a peaceful future. We called this forecast a "conditional forecast," since it is conditional on the assumption that other factors will be negligible.

If we think that other causal factors will tend to *pacifying on net*, then this would imply a greater-than-10% chance of a peaceful future. One possibility we might entertain is that there are other causal factors that have been helping to drive the downward trend in the rate of conflict, in addition to growth, and that these factors will continue to help over the next century as well.

If we think that other causal factors will tend to *aggravating on net*, on the other hand, then this would imply a less-than-10% chance of a peaceful future. One possibility we might entertain is that the downward trend has been substantially driven by the transition from a bipolar to unipolar world order, but that the transition back to a bipolar or multipolar world order will tend to substantially exacerbate the risk of conflict.

Overall, there is no single obvious methodology for adjusting our "conditional forecast" and using it to update our initial outside view forecast. However, overall, it seems fair to say that the conditional forecast does not suggest that we need to make any radical updates to the outside view forecast. It is broadly congruent with the view that a peaceful

future is unlikely but not implausible.

One simple approach is to just average our outside view forecast with this new conditional forecast that focuses on the effect of growth. This leads us to a new overall estimate of a **15% chance** of a peaceful future.[11]

6.6 Appendix: Complications From Uneven Growth

The conditional forecast we produced in Section 6.4.3 — assessing the likelihood that growth will be sufficient to cause the expected rate of conflict to fall by at least 90%, assuming that other changes in the world have no net effect — was built on top of a number of simplifying assumptions. Among these, one of the most consequential may be the assumption that all states will grow at approximately the same rate. This assumption allowed us to focus on a single variable (global GDP-per-capita) rather than having to reflect on the disparate growth rates of different states. However, this assumption is almost certainly false. We know that, historically, most countries have experienced a great deal of growth, but that a small number have hardly grown at all. We also know that some countries with long histories of growth have been begun to show signs of stagnation in recent decades.

It is worth reflecting, then, on how accounting for interstate variation in growth could have influenced our conditional forecast. One basic point is that the effect of growth on the global rate of conflict will tend to be higher if growth is disproportionately concentrated in conflict-prone countries. Another point is that — for sets of countries that are similarly conflict-prone — the effect of growth will tend to be higher if it is spread out between the countries.[12]

[11] This averaging approach is of course fairly unprincipled. However, it is hard to image that a more sophisticated approach could lead to a radically different figure. Since there will be a great deal of room for reasonable disagreement about the overall figure, in any case, differences in the single-digit range are not worth worrying too much about.

[12] If some subset of countries grows more quickly than the other, then this subset will tend to become increasingly peaceful relative to the other countries within the set and account for an increasingly small portion of conflicts. Any further decline in the rate of conflict within this peaceful subset will therefore have an increasingly small impact on the overall rate of conflict.

Overall, it strikes me as far from immediately obvious how future growth will be distributed. It is also therefore far from immediately obvious how we should adjust our forecast to account for uneven growth. Ultimately, for this reason, I have left the conditional forecast as is.

Chapter 7

Conclusion

7.1 Abstract

This chapter briefly summarizes the preceding thesis. It highlights some notable academic contributions, then describes some directions for further work. It closes by advocating for continued attention to "fundamental questions" in International Relations.

7.2 The Thesis So Far

7.2.1 The Central Conclusion

The last chapter closed with a statement that — on the basis of the evidence and considerations that have been explored by this thesis — **a peaceful future should be considered unlikely but not implausible**.

I gave **15%** as specific example of reasonable probability estimate. Since we needed to make a wide range of *judgment calls* to arrive at this number, fairly different numbers could also be considered reasonable. However, it would seemingly be quite difficult to justify much lower number (e.g. 1%) or a much higher number (e.g. 50%). Although the *spectrum of reasonable disagreement* is quite broad, it certainly is not infinitely broad.

7.2.2 How We Arrived at This Conclusion

It is also worth taking a step back, at this point, to remind ourselves of the ground we have covered to arrive at this conclusion.

We started out, in the introductory chapter, by introducing the Peaceful Future Hypothesis. Then, in the literature review chapter, we explored how the debate over the hypothesis has evolved — and also, in some ways, failed to evolve — since late the eighteenth century.

The next chapter — the first "paper chapter" in the thesis — then further clarified the Peaceful Future Hypothesis, defended it as a hypothesis worth studying, and laid out a two-stage methodology for assessing its likelihood. The chapter then walked through the first stage, which involves producing an intermediate forecast purely on the basis of historical trends. One notable finding was there is some evidence that a substantial downward trend in the risk of conflict began in the mid-20th century.

The subsequent two "paper chapters" focused on clarifying the effect of economic growth on the risk of conflict. The underlying motive for this analysis was to put us in a stronger position to reassess and modify our initial forecast, with a better understanding of a *one* especially noteworthy causal factor that could influence the rate of conflict. The key finding was that growth probably has a substantial pacifying effect. However, we also noted some reasons to expect its marginal effect to decline as states become wealthier.

The "coda" chapter then synthesized the key findings from the preceding chapters. It noted that, since the typical state only began to experience substantial economic growth in the mid-20th century, economic growth might be one important factor driving the current (apparent) downward trend in the rate of conflict. The chapter also noted that it is reasonably likely that many states will grow substantially over the next century. There could even conceivably be a major growth acceleration driven by artificial intelligence. This analysis overall bolstered the notion that a peaceful future is at least plausible. However, the chapter argues that a peaceful future still should not be seen as overall *likely*.

7.3 Academic Contributions

This thesis can, to some extent, be read as an extended forecasting exercise. My hope is that the forecast I have ultimately produced — although still preliminary, in many ways, and unavoidably speculative — has intellectual value in its own right. Of course, the chapter "Will the Future Be Peaceful? How Would We Know?" is in part a defense of the intellectual value of this kind of forecast, at least when produced with sufficient epistemological humbleness and care.[1]

Even if the reader is less inclined to give weight to *output* of the forecasting exercise, though, I hope that they will still find value in the exercise itself. As the adage goes: Sometimes it's not the destination, but rather the journey.

Among other academic contributions, this thesis has:

- Clarified and described the intellectual history of a hypothesis that was once central to the field of International Relations (the Peaceful Future Hypothesis)
- Outlined a general framework for evaluating the Peaceful Future Hypothesis and clarified how previous analysis fits into the framework
- Developed a new statistical model that can be used to estimate trends in historical conflict data — and clarified how we should think about what it *means* for there to be a trend in the data, given how irregularly the rate of conflict evolves
- Highlighted underappreciated flaws in certain historical conflict datasets
- Developed a previously neglected explanation for the "puzzle" that wealthy states experience fewer civil wars
- Clarified how "greed" and "grievance" motives for conflict can coexist and interact

[1] In my view, the strongest argument against forecasting hypothesis such as the Peaceful Future Hypothesis is that we do not yet know how viable very *long-term* forecasting is. While forecasting studies have established that at least some people can produce well-calibrated forecasts of near-term political events, there is, naturally, essentially no *direct* evidence that it is tractable to produce well-calibrated forecasts about what the world will be like a century or more out. Maybe, for all we know, the broad forecasting techniques that work well in the near-term do not work nearly as well in the long-term.

- Presented new statistical evidence that, contrary to the conventional view, states become less likely to fight each other as they grow wealthier
- Developed a novel distinction between "scarcity-constrained wars" and "scarcity-driven wars"
- Shown how several arguments about economic causes of peace can be synthesized and grounded in bargaining theory

If nothing else, grappling with the Peaceful Future Hypothesis has turned out to be a good way to grapple with a range of fundamental issues in International Relations. In some ways, as will now be discussed, this thesis's analysis of the hypothesis is still quite preliminary. I imagine that a fully comprehensive effort to grapple with the hypothesis could at least touch on most corners of the field of International Relations.

7.4 Directions for Further Work

A great deal of work remains to be on the Peaceful Future Hypothesis.

For example, in the first paper chapter, the thesis's analysis of historical trends and attempt to produce a trend-based forecast is certainly not entirely comprehensive. Examples of directions for further work include considering additional statistical models and incorporating them into the overall "outside view" forecast (which averages the predictions of different models), considering alternative priors for steps in the forecast that incorporate uninformative Bayesian priors, and — although this would be a great deal more work — working to develop more comprehensive conflict datasets.

The most obvious gap in this thesis, though, is that it considers only a single causal factor that might influence the rate of conflict. I do not seriously consider other factors such as democracy, ethical norms, or the polarity of the international system. These other factors should be explicitly considered and, on the basis of this consideration, the overall forecast of

the likelihood of a peaceful future should be updated. **Overall, I have only taken a first step in incorporating causal considerations into the forecast**.

A broader issue with my forecast is that I do not directly grapple with the possibility that, over the coming century, progress in artificial intelligence will radically and fundamentally change the world. This possibility — which many AI researchers consider likely — came up in a discussion of the possibility that economic growth will substantially accelerate due to widespread automation. However, if we do eventually get to the point where AI systems can do nearly everything that human workers can do, then the rate of economic growth is going to be far from the only thing that's different about the world.

One first-pass way to think about the possibility of radical AI progress is that it gives us a reason to assign more probability to a substantial trend break occurring in the coming century. This update should then filter into our overall forecast in one way or another. Further analysis of the specific effects of AI progress — for instance, its effects on variables such as democracy, international polarity, and the prevalence of information asymmetries — could then influence our expectations about the direction and magnitude of the potential trend break. Appendix B of this thesis actually contains several pages of analysis along these lines, although I have not yet integrated into the thesis's overall forecast.

I am still very unsure about how accounting for the implications of AI *should* shift our overall forecast. My tentative feeling, based on the analysis in Appendix B, is that it should in fact make us *more* optimistic. At the same time, though, reflecting on the potential implications of AI surely tend to broaden the spectrum of reasonable disagreement about the likelihood of a peaceful future. If one person says that accounting for AI should make us substantially more optimistic, while another person says that accounting for AI should make us substantially less optimistic, it strikes me as hard to dismiss either person's view as clearly unreasonable.

Overall, if we want to make further progress on the Peaceful Future Hypothesis, then there are also many directions for further work. We could, of course, go even further by

actively considering the civil war variant of the hypothesis, given that this thesis has focused mostly on interstate war.

7.5 Closing Thoughts: Fundamental Questions in International Relations

Many academic fields — perhaps most academic fields — emerge from a desire to answer *fundamental questions*: questions that have especially broad implications or especially deep intuitive importance. The focus of these fields then tends to shifts over time to more narrow questions.

The founding figures of physics, like Newton and Hooke, were preoccupied with questions that were more fundamental than the ones that are addressed in the typical physics paper today. The founding figures of modern philosophy, like Descartes and Locke, were preoccupied with questions that were more fundamental than the ones that are addressed in the typical philosophy paper today. And so on.

This pattern is natural. Beyond a certain point, opportunities for further progress on the fundamental questions — the ones that motivated the field's founders in the first place — will become very limited. In the positive case, the questions are simply answered satisfactorily. In the negative case, every really worthwhile avenue for evidence-collecting and analysis is explored, and there is simply nothing left to look at or consider that would warrant meaningfully shifting one's views.[2]

International Relations has followed the same pattern. It emerged largely out of a desire to answer fundamental questions, one of which was: "Will the world become more peaceful?" Although modern IR scholars have not fully moved on from these kinds of fundamental questions, it has over time become increasingly rare for IR scholars to dedicate

[2]In the negative case, research papers may continue to *reference* the unresolved fundamental questions — for instance, in their introductions and conclusions — as inspirations for the more narrow questions they are exploring. However, the focus is on ultimately the narrow questions.

direct, sustained analysis to them.

This thesis has, overall, been unusually focused on fundamental questions. This of course includes the question, "Will the world become more peaceful?" But it also includes other questions, like "Why do rich countries almost never have civil wars?" and "What factors have a big influence on the risk of interstate war?"

Even putting aside the specific conclusions I have argued for, I will consider this thesis a success if I have at least convinced the reader that there is *more to say* on the questions I have been examining. This thesis has managed, I believe, to say some new things about old questions. It has also pointed to some new avenues that deserve exploration. I believe it is still too early for the field of International Relations to move on from the kinds of questions that gave birth to it.

Bibliography

Acemoglu, D., Johnson, S., & Robinson, J. A. (2001). The Colonial Origins of Comparative Development: An Empirical Investigation. *American Economic Review*, *91*(5), 1369–1401.

Acemoglu, D., Naidu, S., Restrepo, P., & Robinson, J. A. (2019). Democracy Does Cause Growth. *Journal of Political Economy*, *127*(1), 47–100.

Acemoglu, D., & Robinson, J. A. (2005). *Economic Origins of Dictatorship and Democracy*. Cambridge University Press.

Aghion, P., Jones, B. F., & Jones, C. I. (2017). *Artificial Intelligence and Economic Growth* (Vol. 23928). National Bureau of Economic Research.

Aksu, E. (2008). *Early Notions of Global Governance: Selected Eighteenth-Century Proposals for 'Perpetual Peace' : With Rousseau, Bentham and Kant Unabridged*. University of Wales Press.

Alesina, A., Reich, B., & Riboni, A. (2020). Nation-Building, Nationalism, and Wars. *Journal of Economic Growth*, *25*(4), 381–430.

Allison, G. T., & Zelikow, P. (1999). *Essence of Decision: Explaining the Cuban Missile Crisis* (2nd). Longman.

Allison, P. D. (2009). *Fixed Effects Regression Models*. SAGE Publications.

Anders, T., Fariss, C. J., & Markowitz, J. N. (2020). Bread before Guns or Butter: Introducing Surplus Domestic Product (SDP). *International Studies Quarterly*, *64*(2), 392–405.

Angell, N. (1913). *The Great Illusion: A Study of the Relation of Military Power to National Advantage*. G.P. Putnam's Sons.

Aschenbrenner, L. (2019). *Existential Risk and Growth*.

Assadi, G. (2023). Will Humanity Choose Its Future? *PhilPapers*.

Bahney, B. W., Iyengar, R. K., Johnston, P. B., Jung, D. F., Shapiro, J. N., & Shatz, H. J. (2013). Insurgent Compensation: Evidence from Iraq. *American Economic Review*, *103*(3), 518–522.

Barnhart, J. N., Trager, R. F., Saunders, E. N., & Dafoe, A. (2020). The Suffragist Peace. *International Organization*, *74*(4), 633–670.

Beck, N., & Katz, J. N. (2001). Throwing Out the Baby with the Bath Water: A Comment on Green, Kim, and Yoon. *International Organization*, *55*(2), 487–495.

Bentham, J. (1927). *Jeremy Bentham's Plan for an Universal and Perpetual Peace*. Sweet Maxwell.

Berge, L., Krantz, S., McDermott, G., & Berge, M. L. (2021). Package Fixest.

Bergman, M. (2018). *More Money, More Crime: Prosperity and Rising Crime in Latin America*. Oxford University Press.

Berman, E., Callen, M., Felter, J. H., & Shapiro, J. N. (2011). Do Working Men Rebel? Insurgency and Unemployment in Afghanistan, Iraq, and the Philippines. *Journal of Conflict Resolution*, *55*(4), 496–528.

Berman, N., Couttenier, M., Rohner, D., & Thoenig, M. (2017). This Mine Is Mine! How Minerals Fuel Conflicts in Africa. *American Economic Review*, *107*(6), 1564–1610.

Blair, G., Christensen, D., & Rudkin, A. (2021). Do Commodity Price Shocks Cause Armed Conflict? A Meta-Analysis of Natural Experiments. *American Political Science Review*, *115*(2), 709–716.

Blattman, C., & Miguel, E. (2010). Civil War. *Journal of Economic Literature*, *48*(1), 3–57.

Bloom, N., Jones, C. I., Van Reenen, J., & Webb, M. (2020). Are Ideas Getting Harder to Find? *American Economic Review*, *110*(4), 1104–1144.

Boehmer, C. R., & Sobek, D. (2005). Violent Adolescence: State Development and the Propensity for Militarized Interstate Conflict. *Journal of Peace Research*, *42*(1), 5–26.

Boix, C. (2003). *Democracy and Redistribution*. Cambridge University Press.

Boix, C. (2024). AI and the Economic and Informational Foundations of Democracy. In J. B. Bullock, Y.-C. Chen, J. Himmelreich, V. M. Hudson, A. Korinek, M. M. Young, & B. Zhang (Eds.), *The Oxford Handbook of AI Governance*. Oxford University Press.

Bolt, J., & Van Zanden, J. L. (2014). The Maddison Project: Collaborative Research on Historical National Accounts. *The Economic History Review*, *67*(3), 627–651.

Bolt, J., & van Zanden, J. L. (2020). Maddison Style Estimates of the Evolution of the World Economy. A New 2020 Update.

Bostrom, N. (2014). *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press.

Bostrom, N. (2024). *Deep Utopia: Life and Meaning in a Solved World*. Ideapress Publishing.

Braumoeller, B. F. (2019). *Only the Dead: The Persistence of War in the Modern Age*. Oxford University Press.

Braumoeller, B. F., & Carson, A. (2011). Political Irrelevance, Democracy, and the Limits of Militarized Conflict. *Journal of Conflict Resolution*, *55*(2), 292–320.

Bremer, S. A. (1992). Dangerous Dyads: Conditions Affecting the Likelihood of Interstate War, 1816-1965. *Journal of Conflict Resolution*, *36*(2), 309–341.

Brennen, T., & Magnussen, S. (2023). Lie Detection: What Works? *Current Directions in Psychological Science*, *32*(5), 395–401.

Brooks, S. G. (1999). The Globalization of Production and the Changing Benefits of Conquest. *Journal of Conflict Resolution*, *43*(5), 646–670.

Buckle, H. T. (1861). *History of Civilization in England by Henry Thomas Buckle* (Vol. 1). Parker, Son, Bourn.

Butcher, C. R., & Griffiths, R. D. (2020). States and Their International Relations Since 1816: Introducing Version 2 of the International System (s) Dataset (ISD). *International Interactions*, *46*(2), 291–308.

Buzan, B., & Lawson, G. (2015). *The Global Transformation: History, Modernity and the Making of International Relations* (Vol. 135). Cambridge University Press.

Carlsmith, J. (2020). How Much Computational Power Does It Take to Match the Human Brain. *Open Philanthropy*.

Carr, E. H. (1939). The Twenty Years' Crisis. *New York: Macmillan*.

Carter, D. B., & Signorino, C. S. (2010). Back to the Future: Modeling Time Dependence in Binary Data. *Political Analysis*, *18*(3), 271–292.

Chiappori, P.-a., & Paiella, M. (2011). Relative Risk Aversion Is Constant: Evidence from Panel Data. *Journal of the European Economic Association*, *9*(6), 1021–1052.

Chiba, D., Machain, C. M., & Reed, W. (2014). Major Powers and Militarized Conflict. *Journal of Conflict Resolution*, *58*(6), 976–1002.

Choi, S.-w. (2011). Re-evaluating Capitalist and Democratic Peace Models. *International Studies Quarterly*, *55*(3), 759–769.

Choucri, N. (1974). Forecasting in International Relations: Problems and Prospects. *International Interactions*, *1*(2), 63–86.

Christensen, P., Gillingham, K., & Nordhaus, W. (2018). Uncertainty in Forecasts of Long-Run Economic Growth. *Proceedings of the National Academy of Sciences*, *115*(21), 5409–5414.

Churchland, P. S., & Sejnowski, T. J. (1992). *The Computational Brain*. MIT Press.

Ciccone, A. (2011). Economic Shocks and Civil Conflict: A Comment. *American Economic Journal: Applied Economics*, *3*(4), 215–27.

Cirillo, P., & Taleb, N. N. (2016). The Decline of Violent Conflicts: What Do the Data Really Say? *The Nobel Foundation, Causes of Peace, Forthcoming, NYU Tandon Research Paper*, (2876315).

Clark, G. (2008). *A Farewell to Alms: A Brief Economic History of the World* (Vol. 25). Princeton University Press.

Coe, A. J. (2015). The Modern Economic Peace.

Cohen, D. K. (2018). *Rape during Civil War*. Cornell University Press.

Collier, P., & Hoeffler, A. (2004). Greed and Grievance in Civil War. *Oxford Economic Papers*, *56*(4), 563–595.

Conway, S. (1989). Bentham on Peace and War. *Utilitas*, *1*(1), 82–101.

Cortright, D. (2008). *Peace: A History of Movements and Ideas*. Cambridge University Press.

Costa, D. L., & Kahn, M. E. (2004). Changes in the Value of Life, 1940 to 1980. *Journal of Risk and Uncertainty*, *29*(2), 159–180.

Cottrell, R. C. (2015). Introduction: Sex and Saigon: Gendered Perspectives on the Vietnam War. *Journal of American-East Asian Relations*, *22*(3), 179–186.

Crone, P. (2015). *Pre-Industrial Societies: Anatomy of the Pre-Modern World*. Simon; Schuster.

Cropper, M., Hammitt, J. K., & Robinson, L. A. (2011). Valuing Mortality Risk Reductions: Progress and Challenges. *Annual Review of Resource Economics*, *3*(1), 313–336.

Dafoe, A. (2011). Statistical Critiques of the Democratic Peace: Caveat Emptor. *American Journal of Political Science*, *55*(2), 247–262.

Dafoe, A., Hughes, E., Bachrach, Y., Collins, T., McKee, K. R., Leibo, J. Z., Larson, K., & Graepel, T. (2020). Open Problems in Cooperative AI. *arXiv preprint arXiv:2012.08630*.

Dafoe, A., & Kelsey, N. (2014). Observing the Capitalist Peace: Examining Market-Mediated Signaling and Other Mechanisms. *Journal of Peace Research*, *51*(5), 619–633.

Dal Bo, E., & Dal Bo, P. (2011). Workers, Warriors, and Criminals: Social Conflict in General Equilibrium. *Journal of the European Economic Association*, *9*(4), 646–677.

Davidson, T. (2021). Could Advanced AI Drive Explosive Economic Growth.

Davies, S., Pettersson, T., & Oberg, M. (2023). Organized Violence 1989 to 2022, and the Return of Conflict between States. *Journal of Peace Research*, *60*(4), 691–708.

Deaton, A. (2013). *The Great Escape: Health, Wealth, and the Origins of Inequality*. Princeton University Press.

Delong, J. B. (1998). *Estimating World GDP, One Million BC-Present*.

Department for Science, Innovation, and Technology. (2023). Capabilities and Risks from Frontier AI.

Dixon, J. (2009). Emerging Consensus: Results from the Second Wave of Statistical Studies on Civil War Termination. *Civil Wars*, *11*(2), 121–136.

Djankov, S., & Reynal-querol, M. (2010). Poverty and Civil War: Revisiting the Evidence. *The Review of Economics and Statistics*, *92*(4), 1035–1041.

Drexler, E. (1987). *Engines of Creation: The Coming Era of Nanotechnology*. Anchor.

Drexler, K. E. (2013). *Radical Abundance: How a Revolution in Nanotechnology Will Change Civilization*. Public Affairs.

Dube, O., & Vargas, J. F. (2013). Commodity Price Shocks and Civil Conflict: Evidence from Colombia. *The Review of Economic Studies*, *80*(4), 1384–1421.

Fearnside, P. M. (1998). The Value of Human Life in Global Warming Impacts. *Mitigation and Adaptation Strategies for Global Change*, *3*(1), 83–85.

Fearon, J. D. (1995). Rationalist Explanations for War. *International Organization*, *49*(3), 379–414.

Fearon, J. D. (2008). Economic Development, Insurgency, and Civil War. *Institutions and Economic Performance*, *292*, 328.

Fearon, J. D., & Laitin, D. D. (2003). Ethnicity, Insurgency, and Civil War. *American Political Science Review*, *97*(1), 75–90.

Fitzpatrick, S., Bramley, G., Treanor, M., Blenkinsopp, J., McIntyre, J., Johnsen, S., & McMordie, L. (2023). Destitution in the UK 2023.

Fukuyama, F. (1992). *The End of History and the Last Man*. Free Press.

Gaddis, J. L. (1986). The Long Peace: Elements of Stability in the Postwar International System. *International Security*, *10*(4), 99–142.

Galor, O. (2011). *Unified Growth Theory*. Princeton University Press.

Garfinkel, B. (2020). The Case for Privacy Optimism [Available at: https://benmgarfinkel.blog/2020/03/09/privacy-optimism-2/].

Garfinkel, B. (2021). Is Democracy a Fad? [Available at: https://benmgarfinkel.blog/2021/02/26/is-democracy-a-fad/].

Garfinkel, B. (2024). The Impact of Artificial Intelligence: A Historical Perspective. In J. B. Bullock, Y.-C. Chen, J. Himmelreich, V. M. Hudson, A. Korinek, M. M. Young, & B. Zhang (Eds.), *The Oxford Handbook of AI Governance*. Oxford University Press.

Garfinkel, B., & Dafoe, A. (2019). How Does the Offense-Defense Balance Scale? *Journal of Strategic Studies*, *42*(6), 736–763.

Gartzke, E. (2007). The Capitalist Peace. *American Journal of Political Science*, *51*(1), 166–191.

Gartzke, E. (2015). Interdependence, Development, and Interstate Conflict. *Emerging Trends in the Social and Behavioral Sciences: An Interdisciplinary, Searchable, and Linkable Resource*, 1–13.

Gartzke, E. (2021). Blood and Robots: How Remotely Piloted Vehicles and Related Technologies Affect the Politics of Violence. In *Emerging technologies and international stability* (pp. 113–143). Routledge.

Gartzke, E., Li, Q., & Boehmer, C. (2001). Investing in the Peace: Economic Interdependence and International Conflict. *International Organization*, *55*(2), 391–438.

Gartzke, E., & Lindsay, J. R. (2015). Weaving Tangled Webs: Offense, Defense, and Deception in Cyberspace. *Security Studies*, *24*, 316–348.

Gartzke, E., & Rohner, D. (2010). To Conquer or Compel: War, Peace, and Economic Development. *Institute for Empirical Research in Economics University of Zurich Working Paper*, (511).

Gartzke, E., & Weisiger, A. (2014). Under Construction: Development, Democracy, and Difference as Determinants of Systemic Liberal Peace. *International Studies Quarterly*, *58*(1), 130–145.

Gartzke, E., & Zhang, J. J. (2015). Trade and War. *The Oxford Handbook of the Political Economy of International Trade*, 419.

Gat, A. (2008). *War in Human Civilization*. Oxford University Press.

Gat, A. (2017). *The Causes of War and the Spread of Peace: But Will War Rebound?* Oxford University Press.

Gelman, A., Hwang, J., & Vehtari, A. (2014). Understanding Predictive Information Criteria for Bayesian Models. *Statistics and Computing*, *24*, 997–1016.

Glaser, C. L., & Kaufmann, C. (1998). What Is the Offense-Defense Balance and Can We Measure It? *International Security*, *22*(4), 44–82.

Gleditsch, K. S. (2002). Expanded Trade and GDP Data. *Journal of Conflict Resolution*, *46*(5), 712–724.

Global Database. (n.d.). McDonald's Restaurants Limited. https://uk.globaldatabase.com/company/mcdonald-s-restaurants-limited

Goldstein, J. S. (2012). *Winning the War on War: The Decline of Armed Conflict Worldwide*. Penguin.

Gordon, R. (2017). *The Rise and Fall of American Growth: The U.S. Standard of Living Since the Civil War*. Princeton University Press.

Green, D. P., Kim, S. Y., & Yoon, D. H. (2001). Dirty Pool. *International Organization*, *55*(2), 441–468.

Hall, R. E., & Jones, C. I. (2007). The Value of Life and the Rise in Health Spending. *The Quarterly Journal of Economics*, *122*(1), 39–72.

Hammitt, J. K., Liu, J.-t., & Liu, J.-l. (2000). Survival Is a Luxury Good: The Increasing Value of a Statistical Life.

Hanson, R. (2000). Long-Term Growth as a Sequence of Exponential Modes. *Working manuscript*.

Hanson, R. (2016). *The Age of Em: Work, Love, and Life when Robots Rule the Earth*. Oxford University Press.

Hegre, H., Metternich, N. W., NygÃ¥rd, H. M., & Wucherpfennig, J. (2017). Introduction: Forecasting in Peace Research.

Hegre, H., Oneal, J. R., & Russett, B. (2010). Trade Does Promote Peace: New Simultaneous Estimates of the Reciprocal Effects of Trade and Conflict. *Journal of Peace Research*, *47*(6), 763–774.

Hegre, H., Karlsen, J., Nygard, H. M., Strand, H., & Urdal, H. (2013). Predicting Armed Conflict, 2010–2050. *International Studies Quarterly*, *57*(2), 250–270.

Hegre, H., & Sambanis, N. (2006). Sensitivity Analysis of Empirical Results on Civil War Onset. *Journal of Conflict Resolution*, *50*(4), 508–535.

Heise, L. L., & Kotsadam, A. (2015). Cross-National and Multilevel Correlates of Partner Violence: An Analysis of Data from Population-Based Surveys. *The Lancet Global Health*, *3*(6), e332–e340.

Helliwell, J., Layard, R., & Sachs, J. (2012). World Happiness Report.

Herz, J. H. (1951). *Political Realism and Political Idealism: A Study in Theories and Realities*. University of Chicago Press.

Hoeffler, A. (2011). Greed versus Grievance: A Useful Conceptual Distinction in the Study of Civil War? *Studies in Ethnicity and Nationalism*, *11*(2), 274–284.

Huntington, S. P. (1996). *The Clash of Civilizations and the Remaking of World Order*. Simon Schuster.

Inglehart, R., & Welzel, C. (2005). *Modernization, Cultural Change, and Democracy: The Human Development Sequence*. Cambridge University Press.

Inklaar, R., De Jong, H., Bolt, J., & Van Zanden, J. (2018). *Rebasing Maddison: New Income Comparisons and the Shape of Long-Run Economic Development* (tech. rep.). Groningen Growth and Development Centre, University of Groningen.

Jaques, T. (2006). *Dictionary of Battles and Sieges: A Guide to 8,500 Battles from Antiquity through the Twenty-First Century*. Bloomsbury Publishing USA.

Jervis, R. (1978). Cooperation under the Security Dilemma. *World Politics*, *30*(2), 167–214.

Jones, C., & Vollrath, D. (2013). *Introduction to Economic Growth*. W. W. Norton Company.

Jones, C. I. (2016). Life and Growth. *Journal of Political Economy*, *124*(2), 539–578.

Jones, C. I. (2023). *The Outlook for Long-Term Economic Growth* (tech. rep.). National Bureau of Economic Research.

Jung, J. H. (2016). A Cross-National Analysis of Religion and Attitudes Toward Premarital Sex: Do Economic Contexts Matter? *Sociological Perspectives*, *59*(4), 798–817.

Kant, I. (1983). *Perpetual Peace and Other Essays*. Hackett Publishing.

Kaplan, R. D. (1994). The Coming Anarchy. *The Atlantic Monthly*, *273*(2), 44–76.

Kaplow, J. M., & Gartzke, E. (2021). The Determinants of Uncertainty in International Relations. *International Studies Quarterly*, *65*(2), 306–319.

Kaplow, L. (2005). The Value of a Statistical Life and the Coefficient of Relative Risk Aversion. *Journal of Risk and Uncertainty*, *31*(1), 23–34.

Keegan, J. (1993). *A History of Warfare*. Hutchinson.

Keen, D. (2012). Greed and Grievance in Civil War. *International Affairs*, *88*(4), 757–777.

Kiernan, B. (2017). Viet Nam: A History from Earliest Times to the Present [Section: Alternative Unifications: Rebellion and Restoration, 1771-1859]. Oxford University Press.

Kirk, D. (1996). Demographic Transition Theory. *Population Studies*, *50*(3), 361–387.

LaDou, J., London, L., & Watterson, A. (2018). Occupational Health: A World of False Promises. *Environmental Health*, *17*, 1–8.

Laity, P. (2002). *The British Peace Movement 1870-1914*. Clarendon Press.

Landers, J. (2005). *The Field and the Forge: Population, Production, and Power in the Pre-industrial West*. Oxford University Press.

Laplace, P.-S. (2012). *A Philosophical Essay on Probabilities*. Courier Corporation.

Layard, R., Mayraz, G., & Nickell, S. (2008). The Marginal Utility of Income. *Journal of Public Economics*, *92*(8-9), 1846–1857.

Leacock, S. (1916). Is Permanent Peace Possible? *Maclean's Magazine*, *29*, 7–8, 77–79, 12–13, 92–93.

League to Enforce Peace. (1918). *Win the War for Permanent Peace: Addresses Made at the National Convention of the League to Enforce Peace, in the City of Philadelphia, May 16th and 17th, 1918; Convention Platform and Governors' Declaration*.

Leatherman, J. (2011). *Sexual Violence and Armed Conflict*. Polity.

Levy-garboua, L., Maafi, H., Masclet, D., & Terracol, A. (2012). Risk Aversion and Framing Effects. *Experimental Economics*, *15*(1), 128–144.

Li, X., Xiong, H., Li, X., Wu, X., Zhang, X., Liu, J., Bian, J., & Dou, D. (2022). Interpretable Deep Learning: Interpretation, Interpretability, Trustworthiness, and Beyond. *Knowledge and Information Systems*, *64*(12), 3197–3234.

Lutz, M. (2024). Moral Naturalism. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford encyclopedia of philosophy* (Summer 2024). Metaphysics Research Lab, Stanford University.

Lynn-Jones, S. M. (1995). Offense-Defense Theory and Its Critics. *Security Studies*, *4*(4), 660–691.

MacAskill, W. (2022). *What We Owe the Future*. Basic Books.

Maddison, A. (2007). *Contours of the World Economy 1-2030 AD: Essays in Macro-Economic History*. Oxford University Press.

Maoz, Z., & Abdolali, N. (1989). Regime Types and International Conflict, 1816-1976. *Journal of Conflict Resolution*, *33*(1), 3–35.

Maoz, Z., & Russett, B. (1992). Alliance, Contiguity, Wealth, and Political Stability: Is the Lack of Conflict among Democracies a Statistical Artifact? *International Interactions*, *17*(3), 245–267.

Markwell, D. J. (1986). Sir Alfred Zimmern Revisited: Fifty Years On. *Review of International Studies*, *12*(4), 279–292.

McCloskey, D. (2004). Review of *The Cambridge Economic History of Modern Britain*. *Times Higher Education Supplement*.

McCloskey, D. N. (2016). The Great Enrichment: A Humanistic and Social Scientific Account. *Scandinavian Economic History Review*, *64*(1), 6–18.

McGrath, S. (2010). Moral Realism Without Convergence. *Philosophical Topics*, 59–90.

Mcguirk, E., & Burke, M. (2020). The Economic Origins of Conflict in Africa. *Journal of Political Economy*, *128*(10), 3940–3997.

Mearsheimer, J. (2013). Has Violence Declined in World Politics?: A Discussion of Joshua S. Goldstein's Winning the War on War: The Decline of Armed Conflict Worldwide. *Perspectives on Politics*, *11*(2), 570–573.

Mearsheimer, J. J. (1990). Back to the Future: Instability in Europe After the Cold War. *International Security*, *15*(1), 5–56.

Mellon, J. (2023). Rain, Rain, Go Away: 195 Potential Exclusion-Restriction Violations for Studies Using Weather as an Instrumental Variable. *Available at SSRN 3715610*.

Micale, M. S. (2018). What Pinker Leaves Out. *Historical Reflections/Reflexions Historiques*, *44*(1), 128–139.

Miguel, E., Satyanath, S., & Sergenti, E. (2004). Economic Shocks and Civil Conflict: An Instrumental Variables Approach. *Journal of Political Economy*, *112*(4), 725–753.

Miller, C., & Bakar, K. S. (2023). Conflict Events Worldwide Since 1468bc: Introducing the Historical Conflict Event Dataset. *Journal of Conflict Resolution*, *67*(2-3), 522–554.

Miller, S. V. (2022). peacesciencer: An R Package for Quantitative Peace Science Research. *Conflict Management and Peace Science*, *39*(6), 755–779.

Mokyr, J. (2016). *A Culture of Growth: The Origins of the Modern Economy*. Princeton University Press.

Morris, I. (2014). *War! What Is It Good For? Conflict and the Progress of Civilization from Primates to Robots*. Profile Books Limited.

Morris, I. (2017). *Foragers, Farmers, and Fossil Fuels: How Human Values Evolve*. Princeton University Press.

Mousseau, M. (2000). Market Prosperity, Democratic Consolidation, and Democratic Peace. *Journal of Conflict Resolution*, *44*(4), 472–507.

Mousseau, M. (2009). The Social Market Roots of Democratic Peace. *International Security*, *33*(4), 52–86.

Mousseau, M. (2019). The End of War: How a Robust Marketplace and Liberal Hegemony Are Leading to Perpetual World Peace. *International Security*, *44*(1), 160–196.

Mousseau, M., Hegre, H., & O'neal, J. R. (2003). How the Wealth of Nations Conditions the Liberal Peace. *European Journal of International Relations*, *9*(2), 277–314.

Mueller, J. (1989). *Retreat from Doomsday: The Obsolescence of Major War*. Basic Books.

Mueller, J. (2004). *The Remnants of War*. Cornell University Press.

Naimark, N. M. (1995). *The Russians in Germany: A History of the Soviet Zone of Occupation, 1945-1949*. Harvard University Press.

Neath, A. A., & Cavanaugh, J. E. (2012). The Bayesian Information Criterion: Background, Derivation, and Applications. *Wiley Interdisciplinary Reviews: Computational Statistics*, *4*(2), 199–203.

Nordhaus, W. D. (2021). Are We Approaching an Economic Singularity? Information Technology and the Future of Economic Growth. *American Economic Journal: Macroeconomics*, *13*(1), 299–332.

Oneal, J. R., Oneal, F. H., Maoz, Z., & Russett, B. (1996). The Liberal Peace: Interdependence, Democracy, and International Conflict, 1950-85. *Journal of Peace Research*, *33*(1), 11–28.

Oneal, J. R., & Russett, B. (1999). The Kantian Peace: The Pacific Benefits of Democracy, Interdependence, and International Organizations, 1885 to 1992. *World Politics*, *52*(1), 1–37.

Osterhammel, J. (2014). The Transformation of the World: A Global History of the Nineteenth Century [Section: Peaceful Europe, Wartorn Asia and Africa]. Princeton University Press.

Ouimet, M. (2012). A World of Homicides: The Effect of Economic Development, Income Inequality, and Excess Infant Mortality on the Homicide Rate for 165 Countries in 2010. *Homicide Studies*, *16*(3), 238–258.

Palmer, G., d'Orazio, V., Kenwick, M., & Lane, M. (2015). The MID4 Dataset, 2002–2010: Procedures, Coding Rules and Description. *Conflict Management and Peace Science*, *32*(2), 222–242.

Palmer, G., McManus, R. W., DOrazio, V., Kenwick, M. R., Karstens, M., Bloch, C., Dietrich, N., Kahn, K., Ritter, K., & Soules, M. J. (2022). The MID5 Dataset, 2011 to 2014: Procedures, Coding Rules, and Description. *Conflict Management and Peace Science*, *39*(4), 470–482.

Paravisini, D., Rappoport, V., & Ravina, E. (2017). Risk Aversion and Wealth: Evidence from Person-to-Person Lending Portfolios. *Management Science*, *63*(2), 279–297.

Peace Congress Committee (Ed.). (1849). *Report of the Proceedings of the Second General Peace Congress, Held in Paris on the 22nd, 23rd, and 24th of August, 1849*. Charles Gilpin.

Pinker, S. (2002). *The Blank Slate: The Modern Denial of Human Nature*. Viking.

Pinker, S. (2011). *The Better Angels of Our Nature: The Decline of Violence in History and Its Causes*. Penguin UK.

Pinker, S. (2018). *Enlightenment Now: The Case for Reason, Science, Humanism, and Progress*. Penguin UK.

Pleskac, T. J., Diederich, A., & Wallsten, T. S. (2015). Models of Decision Making under Risk and Uncertainty. Oxford University Press.

Postan, M. M. (1964). The Costs of the Hundred Years' War. *Past & Present*, (27), 34–53.

Powell, R. (2006). War as a Commitment Problem. *International Organization*, *60*(1), 169–203.

PwC Economics. (2013). World in 2050: The BRICs and Beyond.

Ray, D., & Esteban, J. (2017). Conflict and Development. *Annual Review of Economics*, *9*, 263–293.

Ray, J. L. (1989). The Abolition of Slavery and the End of International War. *International Organization*, *43*(3), 405–439.

Ray, J. L., & Russett, B. (1996). The Future as Arbiter of Theoretical Controversies: Predictions, Explanations and the End of the Cold War. *British Journal of Political Science*, *26*(4), 441–470.

Roodman, D. (2020). On the Probability Distribution of Long-Term Changes in the Growth Rate of the Global Economy: An Outside View.

Rosecrance, R. (1986). *Rise of the Trading State*. Basic Books.

Rosenau, W. G. (1997). Every Room Is a New Battle: The Lessons of Modern Urban Warfare. *Studies in Conflict Terrorism*, *20*(4), 371–394.

Rosenstiel, T. (2006, February). Youth and War.

Roser, M., Arriagada, P., Hasell, J., Ritchie, H., & Ortiz-Ospina, E. (2023). Economic Growth [https://ourworldindata.org/economic-growth]. *Our World in Data*.

Russell, B. (1951). The Future of Man. *The Atlantic Monthly*.

Russell, S. J., & Norvig, P. (2010). *Artificial Intelligence: A Modern Approach*. Pearson.

Russett, B., Oneal, J. R., & Davis, D. R. (1998). The Third Leg of the Kantian Tripod for Peace: International Organizations and Militarized Disputes, 1950–85. *International Organization*, *52*(3), 441–467.

Schmidt, B. C. (2002). Anarchy, World Politics and the Birth of a Discipline: American International Relations, Pluralist Theory and the Myth of Interwar Idealism. *International Relations*, *16*(1), 9–31.

Schneider, G. (2017a). Capitalist Peace Theory: A Critical Appraisal. In *Oxford Research Encyclopedia of Politics*.

Schneider, G. (2017b, January). *Economics and Conflict: Moving Beyond Conjectures and Correlations*.

Schneider, G., & Gleditsch, N. P. (2010). The Capitalist Peace: The Origins and Prospects of a Liberal Idea. *International Interactions*, *36*(2), 107–114.

Schneider, G., Gleditsch, N. P., & Carey, S. (2011). Forecasting in International Relations: One Quest, Three Approaches. *Conflict Management and Peace Science*, *28*(1), 5–14.

Signorino, C. S., & Ritter, J. M. (1999). Tau-b or Not Tau-b: Measuring the Similarity of Foreign Policy Positions. *International Studies Quarterly*, *43*(1), 115–144.

Smil, V. (2019). *Growth: From Microorganisms to Megacities*. MIT Press.

Stasavage, D. (2020). *The Decline and Rise of Democracy: A Global History from Antiquity to Today*. Princeton University Press.

Stearns, P. N. (2020). *The Industrial Revolution in World History*. Routledge.

Stepanenko, K., Bailey, R., Hird, K., Mappes, G., Williams, M., Klepanchuk, Y., & Kagan, F. W. (2022, November). Russian Offensive Campaign Assessment, November 2. https://www.understandingwar.org/backgrounder/russian-offensive-campaign-assessment-november-2

Tabatabai, A. (2017). Negotiating the Iran Talks in Tehran: The Iranian Drivers that Shaped the Joint Comprehensive Plan of Action. *The Nonproliferation Review*, *24*(3-4), 225–242.

Taliaferro, S. S. (2015). *Archer's Tale: An Examination of English Archers During the Hundred Years War and Their Impact on Warfare and Society* [Doctoral dissertation, US Army Command and General Staff College].

Tanaka, Y. (2003). *Japan's Comfort Women*. Routledge.

Tang, J., LeBel, A., Jain, S., & Huth, A. G. (2023). Semantic Reconstruction of Continuous Language from Non-Invasive Brain Recordings. *Nature Neuroscience*, *26*(5), 858–866.

Tetlock, P. E. (2017). *Expert Political Judgment*. Princeton University Press.

Tetlock, P. E., & Gardner, D. (2016). *Superforecasting: The Art and Science of Prediction*. Random House.

Thucydides. (2008). *The Landmark Thucydides* (R. B. Strassler, Ed.). Free Press.

Trammell, P., & Korinek, A. (2023). *Economic Growth under Transformative AI* (tech. rep.). National Bureau of Economic Research.

Trask, A., Bluemke, E., Garfinkel, B., Cuervas-Mons, C. G., & Dafoe, A. (2020). Beyond Privacy Trade-Offs with Structured Transparency. *arXiv preprint arXiv:2012.08347*.

Treisman, D. (2020). Economic Development and Democracy: Predispositions and Triggers. *Annual Review of Political Science*, *23*(1), 241–257.

Tuchman, B. W. (1962). *The Guns of August*. Macmillan.

Tuchman, B. W. (1987). *A Distant Mirror: The Calamitous 14th Century*. Random House Trade Paperbacks.

United Nations Population Fund. (2023). World Population Dashboard. https://www.unfpa.org/data/world-population-dashboard

Van Evera, S. (1984). The Cult of the Offensive and the Origins of the First World War. *International Security*, *9*(1), 58–107.

Van Zanden, J. (2009). Introduction. In *The Long Road to the Industrial Revolution* (pp. 1–14). Brill.

Wasserman, L. (2000). Bayesian Model Selection and Model Averaging. *Journal of Mathematical Psychology*, *44*(1), 92–107.

Weakliem, D. L. (1999). A Critique of the Bayesian Information Criterion for Model Selection. *Sociological Methods Research*, *27*(3), 359–397.

Weil, D. (2009). *Economic Growth*. Pearson Addison Wesley.

Wells, H. G. (1914). *The War that will End War*. Duffield Company.

Wohlforth, W. C. (2014). The Stability of a Unipolar World. In *The Realism Reader* (pp. 383–395). Routledge.

Yates, R. D. (2016). Searching for Peace in the Warring States: Philosophical Debates and the Management of Violence in Early China. *Peace in the Ancient World: Concepts and Theories*, 98–121.

Zaidi, W., & Dafoe, A. (2021). International Control of Powerful Technology: Lessons from the Baruch Plan for Nuclear Weapons. *Centre for the Governance of AI*.

Zaidi, W. H. (2011). Aviation Will Either Destroy or Save Our Civilization: Proposals for the International Control of Aviation, 1920 to 45. *Journal of Contemporary History*, *46*(1), 150–178.

Zhang, B., Dreksler, N., Anderljung, M., Kahn, L., Giattino, C., Dafoe, A., & Horowitz, M. C. (2022). Forecasting AI Progress: Evidence from a Survey of Machine Learning Researchers. *arXiv e-prints*, arXiv–2206.

Appendix A

The Risk of Conflict in Our Lifetimes

A.1 Are We All Going to Die?

Braumoeller (2019) closed one chapter of his seminal book with the following reflection:

> *When I sat down to write this conclusion, I briefly considered typing, "We're all going to die," and leaving it at that. I chose to write more, not because that conclusion is too alarmist, but because it's not specific enough.... [I]t's not at all unlikely that another war that would surpass the two World Wars in lethality will happen in your lifetime.*

Braumoeller based this conclusion on a couple of key beliefs. The first belief, which he arrived at by studying the same long-run conflict rate data that we have been studying, is that there has not been a downward trend in the rate of conflict; rather, he believed, the rate of conflict has been fluctuating up and down, without any substantial underlying trend. His second belief was that there has been no change in the probability that any given conflict will escalate to very deadly levels. In other words, in the language we used in Chapter 3, he believed the *Constant Relative Probabilities Hypothesis*: the hypothesis that the relative probabilities of conflicts of different intensity levels are at least roughly constant over time.

This second belief was based on running a statistical test to attempt to identify a change in the distribution of conflict intensity (defined as the fraction of people in the participating countries who are killed) and failing to detect a change.

Together, these two beliefs imply that we should be quite worried about very deadly interstate wars happening in our lifetimes. If conflicts in general are about as likely to occur as they always have been — and if the conflicts that *do* occur are about as likely to become extremely deadly as they always have been — then it follows that very deadly conflicts are about as likely as they always have been.

In this thesis, my analysis has been at least somewhat more optimistic than Braumoeller's. I have been focusing primarily on the conflict trends question and found that we *do* have some reasons to suspect that there is a meaningful (albeit recent) downward trend in the rate of conflict.[1] Given this finding, it seems like we should be at least somewhat less worried about a worse-than-WWII conflict than Braumoeller's analysis would suggest.

However, this thesis has ultimately been focused on a fairly big-picture question that only *indirectly* speaks to the question of whether we are "all going to die." We have focused only on the overall rate of conflict — including both high-intensity conflicts and comparatively unimportant low-intensity conflicts — and have primarily been looking ahead to the far-off period between 2050 and 2300. Barring medical breakthroughs, the majority of this period will occur well after the lifetime of anyone who is alive today.

It is worth, I think, briefly commenting directly on the question of how likely we are to experience catastrophic conflicts during our lifetimes.

[1] Braumoeller noted that conflict rates have moved in an up-then-down pattern over the past two centuries. His view was that the shift from a bipolar world order to unipolar world order explains why the rate of conflict is lower now than it was during the Cold War, but that we should by default expect the rate of conflict to go back up when world order shifts again. However, he did not ultimately attempt to model the statistical process that generated the dataset or directly test for a time trend. As we reported in Section 6.3.1.1, there does actually seem to have been a downward trend both during the bipolar era of the Cold War and during the unipolar post–Cold War period.

A.2 Some Cursory Analysis

There are various ways we can formulate the relevant question. One version we can ask is:

> **Chance of Experiencing a High-Intensity War:** How likely is the average country to experience at least one extremely high-intensity war (e.g. a war where at least 1% of its population dies between 2025 and 2075)?

There are also other related questions we might want to ask. For example, two other questions are:

> **Chance of Great Power War:** What is the chance that there will be at least one great power war between 2025 and 2075?
>
> **Chance of Nuclear War:** What is the chance that there will be at least one full-scale nuclear war between 2025 and 2075?

This last question might be the most consequential one of all. A full-scale nuclear war could kill an unprecedented portion of the world's population, particularly if it triggers a "nuclear winter" during which crops yields substantially decline. If a full-scale nuclear war happens during your lifetime, then there is seemingly a reasonable chance that you personally would die.

Unfortunately, as noted in Chapter 3, it is practically impossible to identify trends in any of these forms of conflict. In the case of great power wars and high-intensity wars, we have enough data points to roughly estimate the average likelihoods of these types of conflict over past two centuries. However, we do not have nearly enough datapoints do be able to tell if the likelihoods have been increasing or declining. In the case of full-scale nuclear war, of course, all we know is that one has never occurred. This observation allows us to infer

that the average likelihood of nuclear war at least cannot have been very high over the past several decades; beyond that, though, it does not tell us just *how* low the likelihood has been, and it certainly does not tell us anything much about trends.

The forecasting problem becomes much easier if we assume that the Constant Relative Probabilities Hypothesis is true. Again, this is the hypothesis that the relative probabilities of conflicts of different types are at least roughly constant over time. If the hypothesis is true, then we can start out by producing a trend-informed estimate of the likelihood of conflict-in-general. (Much of this thesis has been devoted to the problem of producing this kind of estimate.) We can then convert this estimate of the likelihood of conflict-in-general into an estimate of the likelihood of the relevant type of conflict, by (a) estimating the fraction of conflicts that are of the relevant type and (b) multiplying our initial estimate by this fraction.

Unfortunately, the Constant Relative Probabilities Hypothesis may not be true. Although Braumoeller failed to find statistically significant evidence that the distribution of conflict intensity has changed over the past two centuries, it would — as noted in Chapter 3 — be difficult to detect a meaningful-but-not-radical change even if there was one. The error bars on the parameter values are very large and, in practical terms, the test is mainly a test for whether the Constant Relative Probabilities assumption holds between low-intensity and medium-intensity conflicts; there are too few extremely high-intensity conflicts for changes in their underlying probabilities to easily affect the results. Of course, if we are specifically considering (e.g.) nuclear war, then we could not possibly directly detect a change in the portion of wars that become nuclear wars.

There are also theoretical reasons to think that trends for different types of conflict may not be entirely parallel. For example, writing on the "stability-instability paradox" suggests that — in some circumstances — states may become more likely to engage in conflicts if they become more confident that the conflicts will not escalate. This kind of dynamic could, for example, make it possible for the risk of nuclear war and the risk of

low-intensity conflict to actually move in opposite directions.

A more holistic way to structure our analysis, then, might be to take something like the following approach. I'll focus here on the case of great power war.

1) Start out with the assumption that the annual risk of great power war has not changed over time. Estimate this annual risk on the basis of the historical frequency with which great powers have gone to war. Then use this estimated annual risk to produce an initial forecast of the likelihood of great power war between 2025 and 2075, conditional on the assumption that risk has been and will be constant over time.

2) Now attempt to judge the likelihood that the risk of great power war is declining or will decline over time. One way to do this is to directly analyze great power war data, although this is unlikely to reveal much due to the small number of cases. Another way to do this is to ask how likely it is that something like the Constant Relative Probabilities Assumption is true. The Constant Relative Probabilities Assumption can be analyzed partly through statistical tests, such as the one performed by Braumoeller, but also through the consideration of relevant theoretical arguments or more qualitative forms of evidence.

3) Now produce an updated forecast, which is a probability-weighted average of (a) the likelihood of great power war under the assumption that there is no trend in it and (b) the likelihood of great power war under the assumption that the Constant Relative Probabilities Assumption is approximately true.[2]

4) Revise the overall forecast again, by looking at that the specific great powers that currently exist in the world and asking whether there are any strong reasons to think that they are much more or less likely to fight than the forecast suggests. This might involve looking at the frequency with which these great powers have been involved in

[2]Actually, to be comprehensive, we should also explicitly assign weight to scenarios where the Constant Relative Probabilities Assumption is false but where there is still a large (but substantially different) trend in the risk of great power war. For example, we should also assign weight to the hypothesis that the risk of great power war has tended to decline *much more quickly* than the risk of conflict in general.

conflicts in the past (are they unusually conflict-prone or unusually peaceful compared to past great powers?), looking at the potential motives they have to fight each other (are these motives unusually strong or unusually weak?), looking at how strongly they exhibit other traits that tend to predict war or peace (such as distance and economic interdependence), and so on.

I believe that carrying out this approach would be a worthwhile project, to be pursued in future work.

Appendix B

Advanced Artificial Intelligence and the Future of Conflict

B.1 The Possibility of Revolutionary Change

As discussed in Section 6.4.1.4 of this thesis, most AI researchers believe we will eventually have AI systems that — at least in aggregate — can do all of the things that people can do. Surveys also suggest that the typical researcher believes that the necessary capabilities are more likely than not to be developed within this century (Zhang et al., 2022).

There is also no obvious reason to dismiss this hypothesis. Since the human brain is ultimately a physical object that performs computations on data inputs (Churchland and Sejnowski, 1992), there is no *fundamental* physical or metaphysical reason why artificial systems could not perform the same functions. The field of artificial intelligence has also been working toward this goal for less than a century, and, seemingly, up until very recently, has been working with hardware that offers dramatically less computing power than the human brain (Carlsmith, 2020). Furthermore, the field has seen extremely rapid progress over the past decade, towards systems (such as ChatGPT) with quite general and impressive capabilities (Department for Science, Innovation, and Technology, 2023). While it is certainly

possible that the field will eventually hit a wall that it will never get past, it seems like we should assign at least a meaningful probability to extremely advanced AI systems and extremely high levels of automation eventually becoming possible.

So far, we have only discussed one possible implication of progress in AI: the possibility that economic growth will accelerate substantially or even radically (Aghion et al., 2017). However, presumably, if we have AI systems that are sufficient to make human labor redundant, the economic growth rate would be far from the only thing that is different about the world. It would make sense to anticipate many qualitative transformations in the world. We should also anticipate breaks in longstanding social, political, and material trends, since the two previous periods of *revolutionary* economic change — namely, the Industrial Revolution and the Neolithic Revolution — produced many such breaks (Garfinkel, 2024).

One way to try to account for the potentially transformative impact of AI is to simply introduce more uncertainty into our forecast. For example, we can assign a higher probability to a future trend break in the rate of conflict, while trying to make minimal assumptions about whether this break would be toward or away from peace. It is not immediately clear, though, whether "adding more uncertainty" to the forecast should necessarily lead to a higher or lower probability estimate.

The aim of this appendix section is to dig more deeply into the implications of AI progress for the risk of conflict, by briefly cataloging and commenting on several different pathways by which this progress could influence the risk of conflict.[1] For the most part, I will still be talking about the risk of conflict between human-run states that are *using* AI systems as tools. However, at some points, I will also briefly speculate on conflict between AI agents that are not merely fulfilling the preferences of their human users.

Insofar as this section presents a key lesson, I suppose, the lesson is: *AI progress could radically alter the risk of conflict.* Furthermore, my specific view — after weighing up all

[1]The paper whose focus and approach is most similar to this appendix is Gartzke (2021), which explores the potential implications of the complete automation of warfare for civilian casualties and for the attractiveness of conquest.

the different pathways that I consider in this section — is that AI progress is overall likely to *lower* risk of conflict on net. However, this more specific view is also more tentative.[2]

B.2 Potential Influence Pathways

B.2.1 Changes in Fundamental Rationalist Barriers to Peace

As discussed in Section 5.4.1.2, bargaining theory suggests that there are two *especially fundamental* barriers that can prevent even perfectly rational states from identifying and agreeing to peaceful arrangements that would be mutually preferable to war (Fearon, 1995). These two barriers are "information asymmetries" and "commitment problems."

It therefore seems important to consider how AI progress might either reduce or reinforce these barriers. In this subsection, I will first consider its implications for information asymmetries, then its implications for commitment problems, and then, finally, its implications for the level of rationality and efficiency actors exhibit when bargaining. The upshot is that I think that, on the whole, progress in AI is more likely to reduce these barriers than it is to reinforce them.[3]

B.2.1.1 Changes in Information Asymmetries

As has just been noted, information asymmetries are generally believed to increase the risk of conflict (Fearon, 1995). If two states have conflicting beliefs about each other and about the world, they may also have conflicting beliefs about which peaceful arrangements would leave them better off compared to going to war. For example, if two states both believe that they are very likely to achieve an easy victory in a potential war, then they may both be too unwilling to make concessions to avoid this war.

[2]This section has been particularly inspired by various conversations with Allan Dafoe.

[3]See the section "Strong Multilateral Coordination" in Assadi (2023) for a closely related discussion. This section is also well-aligned with the perspective of Dafoe et al. (2020), although their discussion mainly focuses on AI systems whose capabilities are much less dramatic than the ones we are considering here.

I see at least three pathways by which AI could lead to changes in the severity of information asymmetries: (1) by altering the feasibility of intelligence-gathering, (2) by reducing the "lumpiness" of information, and (3) by enabling effective lie detection or some other form of mind reading.

B.2.1.1.1 The Feasibility of Intelligence-Gathering A basic reason that information asymmetries arise is that states often try to hide information about themselves. They often have reasonable worries that information, particularly information about their capabilities and intentions, could be used against them. Ultimately, there are good reasons why most advanced states invest heavily in counter-intelligence.

New innovations could conceivably shift the balance between intelligence and counter-intelligence. For example, on the intelligence side, it is possible that certain advanced AI systems would be extremely efficient and effective at drawing inferences from large volumes of data.[4] It is also possible that, partly due to AI progress, there will be sensors that can more reliably collect sound, image, and electronic signal data without being detected or blocked.

On the other hand, it is possible that AI-enabled tools would actually make intelligence-gathering more difficult. For instance, in the long-run, extremely reliable AI systems for software vulnerability discovery (e.g. as discussed in Garfinkel and Dafoe, 2019) and the removal of human weak-links from digital systems could make signals intelligence much more difficult. In general, it is possible that heavily automated institutions would be much less transparent to each other than ones packed with humans. Perhaps some AI systems could also be especially effective at generating fake data or otherwise managing deception operations that make it difficult to know what is true.

Furthermore, perhaps AI progress will simply result in the world becoming much more complicated and unpredictable, in a way that makes it very difficult to arrived at shared sets

[4]Imagine an intelligence agency where human analysts have largely been replaced by a mixture of specialized AI systems and generally intelligent AI systems that can work at far higher volume, absorb and weigh far more pieces of information at once, and discern patterns in data that an ordinary human could not.

of information and beliefs about the world. For example, whereas it is possible to make a range of fairly reliable inferences about someone's desires and abilities based purely on the fact that they are human, it may not be so easy to make these kinds of inferences about AI agents. In a somewhat similar vein, if AI will cause economic growth and technological progress to accelerate, then — at least until they inevitably slow back down — the high rate of change in the world could simply make it very difficult to keep up and know all of the same things that other actors do.

Overall, it seems quite difficult to say what the net effect of advanced AI systems would be here.

B.2.1.1.2 The Lumpiness of Information One underlying reason that information asymmetries persist, despite being a major risk factor for war and other inefficient outcomes, is that information is *lumpy*: it can be difficult to impart one piece of knowledge without also imparting other pieces of knowledge you would rather keep private (Garfinkel, 2020).

For example, it can be difficult to impart the knowledge "My country is well-prepared to succeed in a war" without also imparting knowledge about capabilities, war plans, and so on that the other country can take advantage of to more successfully prepare. Similarly, it can be difficult to impart the knowledge "My country is not developing biological weapons in these facilities" without also imparting the knowledge "Here are the economically valuable projects my country's biotechnology companies are developing in these facilities." In these cases, countries are seemingly forced to either share the full lump or to share nothing.

The lumpiness of information is the main thing that makes bluffing possible. It *allows* states to credibly claim: "What I am saying is true, but I am not willing to provide substantial evidence, because this would require giving away too much secondary information." If this kind of excuse was not available, then we could safely conclude that any country that fails to provide substantial evidence for its claims is bluffing.

Fortunately, the lumpiness of information is partly contingent on the technology and

institutions that exist. Many different technologies and institutions have arisen, essentially, from efforts to reduce the lumpiness of information and enable what Trask et al. (2020) have called "structured transparency."

To illustrate this point, we can turn to another domain and consider the simple case of bag searches by police. We may want police officers to know whether people are carrying bombs in their bags, but, for privacy reasons, we may not want them to also know what *else* people are carrying in their bags. By default, when it comes to traditional bag searches, the relevant information is unfortunately lumpy. However, if police officers have access to reliable bomb-sniffing dogs, then they can learn whether a bag contains a bomb without learning anything else about what the bag contains. (They can simply have the dog sniff the bag and note whether the dog barks.) As a result, when a reliable bomb-sniffing dog is available, it will be much harder for someone to refuse to demonstrate that their bag is safe — at least without raising strong suspicions in the process.

It seems conceivable that progress in AI could reduce the lumpiness of information, and therefore reduce information asymmetries, if it leads to the development of advanced AI systems that can effectively function like extremely sophisticated sniffing dogs. For instance, as a somewhat idealized case, suppose we could create an AI system that can do the job of a weapons inspector. In principle, human weapons inspectors are meant to receive a high volume of sensitive information (e.g. everything that is going on inside a facility) and then only report out specific non-sensitive conclusions (e.g. that nothing inside the facility indicates a violation of an international agreement). They are meant, in other words, to function like sniffing dogs and just report out exactly the information that needs to be known. The problem with human weapons inspectors, however, is that they are human, have their own loyalties, and therefore in practice often also leak the sensitive secondary information. Advanced AI systems, on the other hand, would not necessarily need to suffer from these limitations. They could conceivably be designed with fairly hard-coded restrictions on what they can report out, and, for additional security, they could be destroyed or have their

memories "wiped" after they report out their key findings.[5]

It could then be much more difficult for states to justify refusing to allow high-stakes claims from being verified, if there exist AI systems that could reliably verify these claims without leaking additional information. If we want to push this idea far enough — and hypothesize the creation of AI verifiers that are known to be, essentially, perfectly reliable and perfectly leak-proof — then some quite radical possibilities become at least conceivable. In principle, for example, we could imagine a verifier system having continual access to an audio and video feed from the office and home of a country's leader, scanning just for sufficiently strong indications that the leader plans to break their word and invade another country.

Insofar as we are already hypothesizing the existence of AI systems that can substitute for human labor, it overall seems fairly likely to me that AI systems will be able to help with the lumpy information problem at least to some degree. The main issue, in my mind, is that it may be hard to convince an adversary that a given AI system actually is reliable and leakproof (even if it is). Some actors may also simply be very unwilling to subject themselves to AI verification, even if this refusal is in some sense irrational.

B.2.1.1.3 The Feasibility of Lie Detection and Mind Reading

Lie Detection and Mind Reading for Humans When people talk about their own intentions and beliefs, it is especially difficult for them to demonstrate that they are telling the truth. In extreme cases, for people who are sufficiently discreet, it is possible to lie without ever doing anything or saying anything to anyone that would reveal the lie.

However, it seems at least possible that progress in AI will eventually enable better lie detection (Brennan, 2023). For example, it is possible that some specialized AI systems will eventually be able to analyze brain activity in a way that allows them to detect when someone is lying, much more reliably than any person could detect lies. It is also possible

[5]Destroying a human weapons inspector is typically regarded as more ethically problematic.

that some AI systems will be able to more directly glean information about the content of people's thoughts. One indication of the potential here is recent work showing that it is possible to use AI, applied to directly to brain activity data, to at least very roughly record the gists of stories that participants silently imagine telling (Tang et al., 2023).

If effective lie detection and mind-reading become possible, then this might make it easier for the leaders of states or other relevant figures to make it clear when they are telling the truth. By the same token, it might become more difficult to bluff effectively about one's beliefs and intentions. The end result might then be a decline in certain information asymmetries, which have historically been extremely difficult to dissolve (Assadi, 2023).

However, of course, it is not clear that effective lie detection and mind-reading will in fact become possible. Even if they do become possible, people will presumably be reticent to subject themselves to these technologies, due to concerns ranging from lack of trust in the technology (especially lack of trust in claims that it is leak-proof), to personal or national pride, to the simple fact that subjecting oneself to mind-reading technology would be quite strange and unsettling. Even if an extremely reliable lie detectors were developed tomorrow, for example, it certainly seems hard to imagine Joe Biden and Xi Jinping strapping themselves into the machines for their next discussion.

Lie Detection and Mind Reading for AI Agents What about if the relevant actors are AI agents, rather than humans?

By default, it seems to me like it should be more tractable to "read the mind" (i.e. discern the beliefs, intentions, and cognitive processes) of an AI system than it is to read the mind of a person. First, it is very difficult to gather granular information about brain activity; the brain is a mushy, messy system that is encased within a skull, and interfering too much with it will risk a person's life. On the other hand, an AI system is a piece of software, and anyone running the software can have complete information about its state at any given point in time.

Unfortunately, the information we have about an AI system's state is typically very difficult to *interpret*. We currently have only a very limited ability to map our detailed under-the-hood information — which, essentially, correspond to long vectors of numbers — onto recognizable concepts (Li et al., 2022). More to the point, we have difficulty using this information to reliably predict how an AI system will behave. However, it is not clear that there are any hard limits here, and the incentives to make progress on interpretability seem fairly strong. Even if AI systems need, to some extent, to be intentionally designed to be more interpretable and predictable, designing them this way could very well be worthwhile.

The upshot here is that, if we imagine future conflicts between AI agents, then it seems possible that they will find it much easier to make credible claims about their beliefs and intentions than humans currently do. This suggests that AI agents systems might be better-able than we are to limit risk-exacerbating information asymmetries about their own beliefs and intentions (Assadi, 2023.)

B.2.1.2 Changes in the Difficulty of Credible Commitments

As discussed above, commitment problems are also often seen as one of the fundamental enablers of conflict (Fearon, 1995). The idea here is: Even if there is some peaceful arrangement that would be mutually preferable to conflict, states may not be able to credibly commit to abide by this arrangement.

Unlike in the domestic sphere, where a national government can be counted on to enforce agreements between private actors, the international sphere exists in a state of *anarchy*: there is no higher power that can be counted on to enforce agreements between states. Therefore, for an agreement to count for something, states need to trust each other to abide by it. Unfortunately, this trust is often hard to come by, because states will often have obvious incentives to renege or renegotiate, particularly if detecting non-compliance is difficult or if circumstances change in a way that produces a shift in relative power.

Reducing information asymmetries could help with some kinds of commitment issues,

for instance by making monitoring (and thus detection of non-compliance) more feasible. However, even if information asymmetries were completely eliminated, some commitment problems would still remain (Powell, 2006). The famous Prisoner's Dilemma scenario is one idealized example of a pure commitment problem: even if both prisoners in the scenario know if and when the other one will defect, they will both still feel pressure to defect. So the problem here is not that the two prisoners lack any information about each other. The problem, instead, is that they cannot *bind* themselves to abide by their initial agreement. In general, rational actors cannot credibly commit to cooperating in (at least non-repeated) Prisoner's Dilemma-type scenarios.

If AI progress could facilitate binding agreements, therefore, then it could conceivably facilitate peace.

B.2.1.2.1 Binding Commitments for Human Institutions Could advanced AI systems help states to make agreements that are more reliably binding? Insofar as we are imagining extremely advanced AI systems that can do anything that people can do, then this at least seems conceivable. One idea here is that certain autonomous AI systems could be designed to reliably punish certain kinds of misbehavior, while also resisting unilateral efforts to deactivate them or interfere with their functioning (Assadi, 2023). For instance, one could imagine an AI system that manages a bank account, which holds a certain amount of money in escrow, or that has the ability to block certain financial transactions. One could even imagine AI systems that are actually empowered to use force.

Stepping back — and as a preface to entering even more speculative territory — people have of course long recognized the potential value of establishing a world government, one that has the power necessary to compel states to follow rules and abide by agreements. Perhaps the single most obvious reason this has never happened, despite serious intellectual interest in the early- and mid-20th centuries, is that it would be extremely difficult for states to *trust* that any world government would actually use its powers in a reliable, fair, and

limited way. Try to imagine, for example, all of the world's national governments giving up their own militaries and trusting some new world government (with a monopoly on force) to manage any disagreements among them; to put the matter lightly, countries like the United States, Russia, and China are probably going to have some *questions* about that new world government.

It seems possible, however, that heavily automated international institutions could fulfill some of the purpose of a world government without running into trust issues that are quite so severe. As in the weapons inspector exampled discussed above, AI systems could be developed to behave in specific, circumscribed ways and could therefore be trusted in ways that humans — who also have their own biases, loyalties, and ineliminable selfish aims — cannot. Another way to put this point is that, whereas human "leviathans" always need to be selected from a pool of dubious candidates, AI-based "leviathans" could be designed to meet the specifications of the people who will be governed by them. This distinction could increase the odds that the worlds' states ultimately do establish leviathans of some kind, which at least partly ameliorate the problem of anarchy.

B.2.1.2.2 Binding Commitments for AI Agents If we again consider conflicts between AI agents, similar considerations apply. The main novel "feature" of the pure AI case is that AI agents may be able to modify themselves so that they can make credible commitments, without requiring external enforcement.

Suppose, as a thought experiment, that it were possible for the two people in the classic Prisoner's Dilemma scenario to modify their own brains, so that they would *become* the sorts of people who place a huge amount of intrinsic value on honoring their own agreements (at least if it seems like the other person will hold up their end of the agreement as well). It would seemingly be in their interest to do this, since it would enable them to enter into genuinely credible cooperation agreements with each other.

Of course, we cannot modify our own psychologies in this way. However, it seems

comparatively more plausible that AI systems could modify themselves — or allow themselves to be modified — in a relatively unconstrained way. This might be another pathway by which AI agents could get past barriers to credible commitments that currently afflict human institutions.[6]

B.2.1.3 Changes in Bargaining Efficiency and in Rationality

Bargaining theory, as articulated by Fearon (1995), identifies barriers that could prevent actors from avoiding mutually undesirable outcomes — even if they both are *perfectly rational* and can bargain with *perfect efficiency*.

Of course, however, no one is perfectly rational. Famously, people often form unreasonable beliefs and often make unreasonable decisions.

Bargaining is also never a perfectly efficient process: identifying and clarifying options, exchanging information and communicating preferences and beliefs, and going back-and-forth on offers can all take an enormous amount of time. To give one concrete example: the bargaining process for the 2015 Iran nuclear agreement, which was to a large extent motivated by a mutual desire to avoid war, took approximately twenty months to negotiate (Tabatabai, 2017).

Even if two actors are in principle capable of identifying and committing to a mutually beneficial arrangement, therefore, they may not always be able to do this sufficiently quickly or at sufficiently low cost. Crisis bargaining is especially problematic in this regard. For example, during the Cuban Missile Crisis, Kennedy and his advisors continued to think up new options, sub-options, and considerations over the entire two-week period; they also continued to gain a better understanding of the overall situation and the Soviet Union's position over the course of the crisis (Allison and Zelikow, 1999). If they had had even less

[6]On the other hand, though, the potential for modification might also create new credibility problems. For example, one can imagine an AI system modifying itself into a cooperative state, while also setting up an external process that will modify it back into an uncooperative state at some future point in time (and also wiping its own memory of the fact that this process is forthcoming). This is not really a variety of problem that arises in ordinary human life.

time to think, compose messages, and receive messages back, before the crisis fully came to its head, then the outcome might have been different. Bargaining inefficiencies, therefore, may have nearly produced a nuclear war.

Fortunately, it seems to me, AI systems may be able to support rational decision-making and efficient bargaining. In principle, it seems possible for advanced AI systems to offer advice that is better than the advice decision-makers would otherwise be getting. (It is already the case, today, that even some very basic systems — such as credit risk assessment software, used in financial contexts — can support rational decision-making.) Advanced AI systems could also simply generate and assess options more *quickly* than people can. (It is also already the case, today, that general-purpose chatbots such as ChatGPT can process large volumes of text, produce lists of high-level ideas, and produce short pieces of analysis much more quickly than people can.) Furthermore, in the same vein, sets of AI systems should presumably be able to generate and process messages to each other much more efficiently than sets of people can. One can imagine AI systems that serve as high-fidelity representatives for humans and drive negotiations forward largely by communicating and negotiating with each other.

Whether AI systems are ultimately supporting bargaining processes between humans, or whether they are directly bargaining with each other as autonomous agents with their own aims, it seems reasonable to expect them to contribute to an increase in rationality and bargaining efficiency. This could in turn reduce the risk of conflict.

B.2.1.4 Summary

To briefly summarize this subsection: Bargaining theory tends to identify information asymmetries and commitment problems as fundamental barriers to peace, even among actors that are perfectly rational and able to bargain perfectly efficiently. The fact that actors are *not* perfectly rational, and *cannot* bargain perfectly efficiently, then tends to further increase the risk of conflict.

It seems that AI progress could reduce each of these barriers to peace. Although there are some considerations that push in the opposite direction, they overall strike me as weaker than the optimistic considerations.

B.2.2 Changes in Economic Conditions

B.2.2.1 Changes in Prosperity

As we have now discussed at length, progress in AI could accelerate economic growth and bring us to a much wealthier world much sooner than we would otherwise expect. Our analysis in Chapter 5 suggested that this change would probably be pacifying on net, although — even given quite radical growth — it would seemingly be insufficient to actually eliminate war. We have reasons to suspect that the marginal effect of growth will decline as the world becomes wealthier, for instance because economic motives will be less relevant to the set of conflicts that do continue to occur.[7]

B.2.2.2 Changes in the Returns on Consumption

The Diminishing Returns Mechanism that may link growth to peace depends on people experiencing sharply — specifically, *sub-logarithmically* — diminishing returns on consumption. (See Section 5.4.1.3 for a discussion.) If people came to experience less sharply diminishing returns, then the standard direction of the effect actually reverses. Since growth increases the amount of wealth to fight over, people who do not experience sharply diminishing returns should actually — all else equal — become more motivated to fight over it. The new logic would be: the larger the potential prize is, the more important it is to hold a large portion of it.

[7]Stepping back, and looking even further ahead, it may also be worth noting that even enormous increases in prosperity (at the per-capita level) may be temporary in the very long run (Bostrom, 2024). There are presumably physical limits to how much economic output can be produced within a limited volume of space. That means that total output must eventually hit a ceiling. If it does, then prosperity will last only if population levels do not eventually rise high enough to bring output-per-capita back down to a low level. Ultimately, the Malthusian dynamics that preceded the modern emergence of rapid economic growth could return.

If the returns on people's consumption start to decline less sharply, therefore, then prosperity could actually become a *risk factor* for conflict.

B.2.2.2.1 Changes Driven by the Introduction of New Products Why might the returns on consumption change? The most obvious explanations involve changes in the goods and services that are available to people.

First, returns curves probably become less sharply diminishing as product variety increases. To illustrate this point, we can imagine an economy where water, potatoes, and tents are the only items for sale. Life on a $500 income might be only somewhat harsher in this economy than it is in ours. However, life as a millionaire would be far worse. A million dollars worth of water, potatoes, and tents is simply far less enticing than a million dollars worth of the basket of goods available to millionaires today. Seemingly, the returns on consumption diminish more quickly when there is less product variety.

If technological progress continues, then we should expect the range of goods and services available to continue to increase. Particularly since progress in artificial intelligence may lead the rate of innovation to accelerate, however, we may be moving toward a world with dramatically more things to buy. This shift could at least conceivably cause the returns on consumption to become superlogarithmic rather than sublogarithmic.

It is also at least conceivable that people's demand for some specific new products, enabled by progress in AI, will be particularly insatiable. I suspect that some drug addicts, for example, do not in practice behave as though there are sharply diminishing returns on their drug of choice (at least up to a point). Perhaps there will be some new products that the typical person can continue to spend more and more on, without the marginal value dropping nearly as sharply as it does with other products. Bostrom (2024) gives the illustrative example of extremely expensive but genuinely effective life extension products.

B.2.2.2.2 Changes Pertaining to the Emergence of AI Agents Perhaps, as also raised above, some AI systems will behave like autonomous actors in their own rights, with

objectives that do not simply mirror those of human users or overseers. This might happen either because people decide to build these kinds of AI systems, for one reason or another, or because there is at least some degree of unintended loss of control (Russell and Norvig, 2010; Bostrom, 2014; Department for Science, Innovation, and Technology, 2023).

I see no obvious reason to assume that autonomous AI agents will behave as though they have the same returns curves that people do. The rate at which people become satiated presumably has a great deal to do with contingent aspects of human psychology and biology, which have been shaped by our idiosyncratic evolutionary history. So AI systems could be different.

Extremely speculatively, at least if we are imagining a world where AI agents are freely competing with each other for resources, it seems like we might actually expect them to satiate less quickly than people do. If we draw a rough analogy between what food means for people and what computational resources means for AI systems, then — when it comes to survival and self-propagation — the "inherent returns" on food do seem to be far shallower. Beyond a certain point, in fact, eating more is actually harmful for a person's survival and reproductive success. More and more computing power, however, could for example be put to use running more and more instances of an AI agent. If the utility functions of AI agents will at least eventually evolve through a kind of Darwinian process, with some AI agents gaining more access to resources and thereby becoming more prevalent or influential, then it seems like the systems that emerge from this process would display less satiation behavior than humans currently do. Therefore, it is at least conceivable that prosperity would actually be a risk factor for conflict between AI agents.

On the other hand, it is also at least conceivable that AI systems would exhibit stronger satiation behavior than people do. For example, they might be intentionally designed this way, specifically to limit their inclination to take risky actions. It is therefore also possible that prosperity would have an even stronger pacifying effect on AI agents than it does on people.

B.2.2.3 Changes in Economic Interdependence and the Seizability of Wealth

The strange kind of world we are imagining, where extremely advanced technology exists and human labor plays only a minimal role in economic production, is also presumably one where the economy is fairly structurally different from our current economy.

Should we expect this economy to exhibit more or less economic interdependence? One way to get a handle on this question is to ask whether there are any reasons to expect manufacturing supply chains to become more or less complex and globally distributed. A basic argument in favor of "more complex" is just that some future technologies will presumably be even more complex than those that exist right now; complex technologies, which integrate many different highly specialized parts, also tend to have more complex supply chains. In general, so far, the clear long-run pattern has been that economic growth and increasingly complex economic networks go hand-in-hand.

On the other hand, one argument for "less complex" is that technological progress may also enable new approaches to manufacturing. In particular, suppose that 3D-printing-type manufacturing processes become increasingly sophisticated over time; in the physical limit, this could conceivably involve "atomically precise manufacturing" processes that make it possible to essentially print specific objects atom-by-atom, using only common chemical elements (Drexler, 2013). This more integrated approach to manufacturing, using raw material that is available anywhere, would presumably involve less distributed supply chains and therefore make states less dependent on each other (and therefore less averse to disrupting each other's economies through conflict). However, the ultimate feasibility of this kind of manufacturing remains unclear.

Regarding the nature of wealth: it seems fairly natural to imagine that computing hardware and energy sources (to run this hardware) would become much more important as sources of wealth. If one state could capture some of another state's datacenters and adjacent power plants, without physically damaging them, then they might serve as attractive prizes. Intuitively, they might also be more feasible to capture than at least some other modern

sources of wealth, particularly if their operation and maintenance is automated and does not rely on staff with rare tacit knowledge. Beyond this, I am not sure how to imagine exactly how else wealth would manifest in this radically different economy.

Gartzke (2021) makes another, related point, which also bears on the seizability of wealth: regardless of what form wealth takes, autonomous systems may generally be capable of seizing wealth more efficiently (i.e. with lower operational costs and with less unintended destruction of value) than human combatants currently can. This would then tend, all else equal, to make conflict more attractive.

Overall, it strikes me as particularly unclear whether structural economic changes brought by AI-driven automation and technological progress would support or undermine peace on net.

B.2.3 Changes in the Prevalence of Democracy

Up until the past two centuries, virtually all large and powerful governments were autocratic (Stasavage, 2020). As discussed in Section 5.5.3, given how strongly democracy and economic development have coincided, it seems plausible that changes in economic conditions have at least *facilitated* the modern emergence and spread of democracy.

If that is right, however, then it also suggests that *further* changes in economic and technological conditions could conceivably undermine democracy. Even if global conditions are currently relatively ripe for democracy (at least compared to the conditions in previous centuries), they may not be that way forever.

There are a number of possible stories we could tell about how highly advanced AI could undermine democracy on a large scale (Garfinkel, 2021; Boix, 2024). The most obvious, in my view, is that it might be bad for democracy if (a) human labor becomes largely worthless and (b) surveillance and law enforcement can largely be automated. In this world, as Gartzke (2021) notes, people who do not own substantial capital have essentially nothing they can threaten to withhold (insofar as their labor is not needed) and very little they can

threaten to do (insofar as automated law enforcement is highly effective). These same people might also be naturally inclined to demand heavy redistribution, insofar as they would be dependent on government welfare or private charity, and insofar as the default distribution of wealth would obviously have nothing to do with merit.[8] Together, these considerations suggest that wealthy elites would have weak incentives to devolve their power, while also having good reasons to expect power devolution to be financially costly to them. This does not seem like a recipe for robust democracy.

That being said, it is also possible to imagine progress in AI supporting democracy. For example, if the dynamics of diminishing strongly returns cause people to care less about the exact division of wealth than they once did, then elites may become less bothered by the redistributive implications of democracy than they would be in poorer societies. It is also conceivable that AI could support novel forms of democratic participation (e.g. via AI systems that accurately represent individuals' views in deliberations) or that applications of AI could in some way be used to protect against democratic backsliding (e.g. by helping to increase government transparency, making it possible to detect when a political figure has anti-democratic intentions, or by automating the enforcement of constitutional constraints).

AI progress could also influence the prevalence of democracy by reshuffling the global balance of power, since some states may harness it much more successfully or quickly than others do. For instance, it seems possible that the future will be more democratic if the United States is massively globally dominant than it would be if China is massively globally dominant. As an analogy, the fact that Britain industrialized early and led the world for a lengthy period of time probably had flow-through effects on many states' political institutions in the 20th century.

Since democracies are widely believed to be less likely to fight each other, while democracy-autocracy pairs may actually be even more likely to fight each other than autocracy-autocracy pairs are, any sufficiently large change in the prevalence of democracy

[8]Similar to the way things were in the pre-industrial era, when land was the major source of wealth (Crone, 2015), wealth accumulation would essentially be synonymous with rent collection.

could have a meaningful effect on the risk of conflict (Gartzke and Weisiger, 2014). That being said, given that we have only a very poor understanding of how regime type influences conflict, it is also possible that the causal mechanism — whatever it is — would not actually be nearly as applicable in a world radically transformed by AI.

B.2.4 Changes in the Relationship Between Offense and Defense

B.2.4.1 Changes in the Offense-Defense Balance

In most cases, the easier it is to carry out some kind of attack, the more likely this kind of attack is to happen. For example, it is easier to invade a country by crossing a land border than it is to invade a country by crossing an enormous body of water. Therefore, countries tend to invade their direct neighbors much more often than they invade countries that are separated by water. However, because good ships make it easier to invade across water, the development of good ships in the modern era has increased the frequency of invasions taking place across water.

Relative to a particular kind of attack, and relative to a particular pair of actors, an *offense-defense balance* describes how easy it is for the first actor to carry out the attack if the second actor also puts some amount of effort into defending itself (Garfinkel and Dafoe, 2019). *Offense-defense theory* is then a body of International Relations literature that primarily grapples with the implications of shifts in offense-defense balances (e.g. Jervis, 1978; Lynn-Jones, 1995; Glaser, 1998).

If an offense-defense balance tips toward offense, or at least is perceived to have tipped toward offense, then the relevant attack should *typically* become more likely. However, deterrence dynamics prevent this pattern from being *universally* true. In particular, if some form of attack becomes extremely viable for both sides of a dyad, then a corresponding increase in the ease of retaliation may lead the two states to both become more reticent to carry the attack out. This is, in fact, the situation we observe in the context of nuclear strikes and missile defense.

Presumably, progress in AI will enable new military technologies and new doctrines that influence offense-defense balances to some degree. It seems, a priori, difficult to predict what the overall direction of this effect will be. Military analysts have a very mixed track record when it comes to forecasting offense-defense balances. Famously, in cases like the First World War, analysts have often failed to correctly evaluate the balances that exist even at particular points in time (Van Evera, 1984).

Garfinkel and Dafoe (2019) suggest that it may be easier to interpret the implications of *quantitative* technological changes than it is to interpret the implications of *qualitative* changes. For instance, suppose we want to predict the implications of highly effective software vulnerability discovery tools. This may be relatively tractable, since their effect may be similar to the effect of simply scaling up the kind of software vulnerability discovery work that is already being done by both attackers and defenders.

We suggest that, in a many cases, it may make sense to expect to observe a pattern of *offensive-then-defensive scaling* ("OD-scaling"): this means that scaling initially tilts the balance toward offense, but in the long-run tilts the balance toward defense. The underlying dynamic here is that scaling often initially helps the attacker to exploit neglected spots in the other party's defenses (e.g. software vulnerabilities the defender has not yet discovered and patched), but continued scaling can eventually allow the defender to cover up all of their key weak spots (e.g. by discovering and patching all of their consequential software vulnerabilities) and thereby make it difficult for the attacker to find a way through. However, this pattern is far from universal. We present missile and drone defense as a counter-example, where it seems comparatively difficult to imagine a defender covering up all of the "weak spots" in their national airspace; adding more offensive missiles and more defensive missiles to both sides will probably just continue to help the attacker over and increase the odds that at least some of their offensive missiles get through.

What this kind of offense-defense balance analysis will add up to, in the end, is far from clear. Even if we can get a handle on the implications of quantitative shifts, the

qualitative ones seem extremely hard to reliably foresee. One *further* forecasting difficulty is that offense-defense balances are in fact interrelated. For example, if it becomes easier to conduct cyberattacks that disrupt communications systems, then it may also become harder to conduct physical invasions that rely on reliable access to communications systems (Gartzke and Lindsay, 2015). Therefore, even if we can confidently predict that some technology change will increase the ease of a specific form of attack, this does not necessarily imply that the overall risk of large-scale war will increase.

In my mind, it would currently be hard to justify substantially revising one's predictions about the future of conflict on the basis of offense-defense balance considerations. That being said, I believe that the most interesting question to explore further is whether we should expect at least some possible attacks with weapons of mass destruction — particularly attacks with nuclear weapons, but also attacks with "substitute WMDs" such as biological weapons — to continue to be both very difficult to defend against and very difficult to prevent with preemptive strikes. The answer to this question would seem to have especially clear implications for both the likelihood of major conflict (due to deterrence effects) and the potential stakes of major conflict.

B.2.4.2 Changes in Offense-Defense Distinguishability

Within the offense-defense theory literature, scholars also sometime discuss the concept of offense-defense balance distinguishability (Jervis, 1978). *Offense-defense distinguishability* describes the extent to which it is possible to make primarily defensive military investments, as opposed to military investments that facilitate both offensive and defensive operations to similar degrees.

For example, buying firearms and tanks will tend to be very useful both for defending your own territory and for invading someone else's territory. On the other hand, bunkers built on one's own territory are of course more useful for defense. If the most defensively-useful investment portfolios and the most offensively-useful investment portfolios are nearly

the same, then that means that "offense-defense distinguishability" is low.

The concern here is that, if offense-defense distinguishability is low, then states may end up engaging in costly arms build-ups that ultimately fail to increase anyone's security. Consider, for example, a pair of states where both states *mostly* care about lowering their own chances of being invaded. The first state, nervous about this risk, may make some military investments to enhance its own ability to defend itself. Unfortunately, because these investments are also very useful for offense, the investments have the knock-on effect of reducing the second state's security. (The second state may also begin to wonder whether the investments indicate that the first state is *preparing* to invade it.) The second state will therefore respond by making some additional defensively-motivated investments of its own. Unfortunately, these additional investments once again undermine the first state's security. Perhaps the first state will then want to make some further investments....

If offense-defense distinguishability is high, on the other hand, then the first state can make investments that are mainly useful for defense.[9] These investments will have a smaller effect on the security of the second state and therefore trigger a smaller investment response on its part. This second state's (smaller) investment will also probably be primarily tilted toward defense, which will further limits its negative effect on the first state's security. The ultimate results are (a) the two states invest less overall and (b) the investments that the two states make actually *do* end up having a net negative effect on the viability of invasions. For our purposes, the latter result is the one that is particularly of note.

Drexler (1987) argues that progress in AI could ultimately enable greater distinguishability. In particular, he argues that we should eventually be able to create tamper-proof weapons systems that are programmed to use force only if certain conditions are met. He gives the example of an space-based weapon which has been programmed only to operate if

[9]No investment is ever *purely* defensive, since, at minimum, any good defense will tend to reduce the other state's ability to deter future attacks through credible threats of retaliation or weaken ongoing attacks by damaging the attacker's power base. As a famous example of an investment that might seem purely defensive at face-value, but was not purely defensive in practice: the long-wall built by Athens ahead of the Peloponnesian War was (correctly) perceived by Sparta as decreasing its security (Thucydides, 2008).

its sensors indicate that the territory of the country that launched it is under attack. More broadly: Insofar as we are assuming that AI systems will eventually be able to substitute for virtually all human labor, it then seems likely that, beyond a certain point, virtually all weapons systems will be autonomous weapons systems; it should also be possible for the software managing these systems to make very sophisticated judgements. Although the idea is somewhat whimsical, we can imagine an autonomous tank that will refuse to move unless it is convinced that the current war was genuinely provoked by the other state. Drexler uses the term "active shields" to refer to autonomous weapons systems that have these kinds of restrictions.

The feasibility of this kind of proposal will ultimately depend, though, on at least: (a) whether autonomous systems can be made tamper-proof, so that their programming cannot be modified and so that their physical components cannot be removed without destroying them, and (b) whether the party acquiring the weapons system can *convince* other parties that the system is *genuinely* defensively programmed and tamper-proof. If these conditions are met, then AI progress could conceivably reduce the risk of conflict in part by increasing offense-defense distinguishability.[10]

B.2.5 Changes in International Order

The distribution of power in the international system probably influences the risk of conflict (e.g. Braumoeller, 2019). For example, unipolar international systems may tend to have lower rates of conflict than bipolar ones, because the hegemonic state can enforce peace between the other states and does not feel pressure to engage in rivalries or proxy wars. In the limit, something resembling a full-on world government would presumably exert the same pacifying effect to an even greater degree.

AI progress could conceivably lead to a more strongly unipolar world order. This

[10]There may also be pathways by which AI progress will influence offense-defense distinguishability. Conceivably, if future military systems will be very different among many dimensions, then the extent to which they exhibit offense-defense distinguishability could also be very different, in one direction or another.

might happen, for example, if the AI industry exhibits strong winner-take-all dynamics (e.g. more users -> more data -> better AI systems -> more users) and eventually underlies essentially all economic output. If AI-driven growth will be fast and particularly concentrated within the leading AI power at the time, that could also lead a single state to rapidly move far out ahead of others; perhaps the state could then find ways to lock in this advantage for a long time. This would be analogous to the way in which the Industrial Revolution substantially boosted Britain above its rivals for a time, even though its advantage did ultimately erode as industrialization spread.[11]

At the same time, AI progress could also at least conceivably lead to essentially any kind of world order, depending on the extent to which different states successfully take advantage of it. Although I see no *specific* reason to think that a bipolar world order would become more likely, for example, it is also certainly possible to imagine two major states leveraging the technology particularly well, without either pulling dramatically far ahead of the other.

[11]The most extreme version of this possibility is one where the leading state becomes so dominant that it essentially becomes a *de facto* world government, imposing its will on other states while facing virtually no risk of successful resistance (Bostrom, 2014; Assadi, 2023). Should we think this is plausible?

To my mind, the three key variables here are: (a) how much more wealthy and technologically advanced the leading state would become relative to the lagging states, (b) how effectively its economic and technological lead could be *converted* into the power to coerce or compel the other states, and (c) how motivated the leading state actually would be to apply this power and accept any risks involved in applying it. Interestingly, even if one country does end up jumping out unprecedentedly far ahead of all others, the latter two variables could still block anything akin to world conquest from occurring.

Let's consider the second variable. As an analogy, even though South Korea's GDP is about forty times North Korea's GDP, South Korea has only a very limited ability to impose its will on North Korea: North Korea's ability to deter South Korea, by threatening the use of a wide array of conventional and non-conventional weapons that South Korea cannot easily defend against, remains strong. The point here is that if the "offense-defense balance" for at least one form of high-severity attack (e.g. attacks with nuclear weapons) remains or becomes sufficiently strongly tilted toward offense, then deterrence could end up limiting the ability of even extraordinarily dominant states to coerce or compel other states.

Now, to comment on the final variable: even if something akin to world conquest does become viable, at least within some window of opportunity, this certainly does not guarantee that the leading state would actually *attempt* it. For instance, suppose that, at some future point in time, the US will have capabilities that it believes may be sufficient to neutralize the nuclear arsenals of other states and remove their abilities to pose military threats: the level of *confidence* that the US would need to have in its newfound abilities, in order to actually try putting them to use, would need to be extraordinarily high. The fact that the application of this neutralizing power would tend to violate norms would also, I imagine, be a genuinely important consideration. Countries like the US already frequently refrain from using or threatening to use force to get their way, even in cases where they in principle could.

If we wanted to construct a case for a highly multipolar world, then perhaps the case would be: beyond a certain point, it will be hard for states to maintain advantages over each other, because economic production will be essentially fully automated and will be underpinned by essentially the same AI software everywhere. (At some point, the software developed by leading countries is bound to leak or be redeveloped.) Growth and innovation will also at some point need to approach some kind of ceiling, due to physical limits, preventing the leading states from continuing to run forward to stay ahead. In this world, then, maybe many states would be similarly powerful. IR theory then suggests that, in such a multipolar world, the risk of conflict would by default by higher.

If we again imagine conflicts between AI agents or collections of AI agents, rather than conflicts between states, the polarity issue seems to become especially difficult to think about. For now, then, I will simply not think about it.

B.2.6 Changes in Culture and Values

Finally, in a world radically transformed by AI, human culture would presumably be different in some way. I am not sure what this way would be. However, I will gladly engage in one last burst of speculation.

B.2.6.1 The Direction of Change

As a first pass, maybe we should expect current trends in human values — which tend to coincide with economic development — to carry forward or even to accelerate. This might suggest, for example, cultures that are even more individualistic, tolerant, feminized, self-expression-focused, rule-following, and violence-averse. (See Section 5.5.4.) I expect that further change in this direction would, overall, tend to make states less likely to fight.

However, particularly since we do not understand why economic development tends to produce these changes (if in fact it does), it seems hard to be *confident* that the pattern would continue to hold even as the circumstances people are experiencing become radically

different. For example, if paid work ceases to become a part of human life, then that would surely have some distinctive cultural implications.

How about culture and values among possible AI agents? Now, of course, we are *really* speculating. Insofar as these agents can be understood to have values, it seems like there are a few possible origins for these values. First, the AI agents could (at least in part) have values that are derived from those of humans. Second, the AI agents could (at least in part) have values that emerge through evolutionary processes and thus are selected for being competitively advantageous. Third, the AI agents could (at least in part) have values that emerge somewhat randomly or idiosyncratically.

In the first case, we are back to the question of what human values will be. In the second case, we can try to reason about what values would be competitively advantageous, in a hypothesized world where AI agents need to compete with each other for resources. Although Hanson (2016) does make some educated guesses, this case is also of course extremely difficult to reason about. The third case might be the most difficult of all. We know that individual humans have ended up with plenty of idiosyncratic values that would have been very difficult to foresee. It does not seem like it would have been easy to anticipate, on the basis of our evolutionary history, that some humans would end up deciding to dedicate much of their free time to getting better at unicycling, or deciding to shoot a golden phonograph record out into the depths of space, or deciding to burn themselves alive in order to advocate for the protection of Vietnamese Buddhists.

One generally interesting question — both for people and for AI agents — is how strongly we should expect future actors to be motivated by *ethical values*, in particular values that concern aspects of the world beyond their own lives and experiences. Anti-violence and anti-coercion values would presumably be pacifying. However, beyond these values, it may actually be bad news for the future of conflict if actors become increasingly ethically motivated — at least if their ethical values are not sufficiently aligned with each other (Braumoeller, 2019). We saw in Section 5.4 of this thesis that sufficiently ideologically

motivated actors may actually be willing to suffer significant economic losses in order to fight for the causes they care about. More broadly, the more actors care about the world beyond themselves, the more potential there is for conflicts between their preferences. If I don't just care about what happens inside my own home, but I also care about what happens inside *your* home, then there will naturally be a greater chance of an important discrepancy arising between my aims and yours. Speculatively, actors primarily pursuing ethical values may also be less prone to satiation when pursuing the accumulation of resources (Bostrom, 2024): if someone is mostly hoping to use a pool of money to save lives, for example, rather than hoping to use the money to buy themselves nice things, then perhaps they will tend to assign the money a value that is closer to linear. Insofar as this is true, this would be another reason why more ethically-minded actors could sometimes be more willing to engage in conflict.

Ultimately, it strikes me as very unclear in what directions culture and values will change. Although one can imagine them changing in a pacifying direction, it seems about as easy to imagine them changing in a direction that makes conflict more likely.

B.2.6.2 Should We Expect Convergence in Values?

On the other hand, optimistically, maybe the values of different actors will tend to *converge* in a way that reduces the risk of conflict between them (Assadi, 2023). As an idealized case: if everyone cared more about their ethical aims than they cared about their own personal interests, and if everyone had roughly the same ethical aims, then they would have very little reason to fight.

Why might people converge in their ethical aims? One possible story is that some moral views are *recognizably correct* and that most people would converge on these correct moral views if only they reflected on them carefully enough or had the arguments for these views explained well enough. Future AI systems might help people do the necessary moral reflection more efficiently and reliably, at least if people choose to use these AI systems for

this purpose. People might then increasingly converge in their moral views.

I am, personally, inclined to be skeptical of the above story. There are practical and psychological grounds for skepticism, but philosophical ones as well. First, it might simply be false that any particular set of moral views is genuinely "correct." Furthermore, some philosophers (e.g. McGrath, 2010) have argued: even if there *are* genuine moral facts, we should not *necessarily* expect even fully rational and fully informed actors to converge upon them.

Here's another story we could tell: maybe the differences between people's fundamental moral beliefs are actually much smaller and much less consequential than they seem, because moral disputes actually mostly flow from differences in people's beliefs about non-moral claims (Assadi, 2023).[12] As an illustrative example, consider the conflict between Israel and Hamas. This conflict is clearly - at least to some degree - morally motivated. At the same, though, most people on the two sides of the conflict have genuinely different views about historical events; about the motives and ambitions of particular people and groups; about whether and why particular tragedies occurred; and about God, the afterlife, and so on. These differences in beliefs about non-moral claims seem to play a role in explaining why the two sides have such divergent beliefs about which actions would be morally justified and about which outcomes would be morally good. If differences in beliefs about non-moral claims *do* play an important role in driving morally motivated conflicts, then perhaps that should make us more optimistic. Some scholars argue that people have, on average, been becoming gradually more rational and well-informed over time and thus more capable of forming accurate beliefs (Pinker, 2018). If that is so, then perhaps further technological progress (e.g. in AI systems that provide people with high-quality information and help them think through problems) could lead people to converge more strongly on accurate non-moral beliefs. Convergence on these non-moral beliefs might then, to a large extent, support convergence in people's morally motivated aims.

[12]Note, though, that some philosophers — specifically, philosophers subscribing to "analytic naturalism" — deny that there is a fundamental distinction between moral and non-moral claims (Lutz, 2024).

On the other hand, cynically, maybe people are simply too biased or too unworried about holding false beliefs for the availability of even extremely reliable information sources and extremely useful thinking tools to make much of a positive difference. The absence of reliable information sources and useful thinking tools is not obviously the biggest bottleneck preventing, for example, Democrats and Republicans in the US from having already converged on the same views.

Furthermore, it is also certainly possible to imagine people's beliefs diverging more than they do now. For instance, maybe we should worry about governments leveraging AI to more effectively control people's information environments and shape their beliefs.

How about AI agents, though? Would we have any reason to expect them to converge in the ethical aims they pursue? One possible story is that most AI agents could share similar aims simply *by design*: maybe they would be created by only a handful of human actors, who mostly share similar views. Another story is that, if AI agents ultimately evolve through a competitive process, then the AI agents that emerge from this process could conceivably be copies or close variants of a particular AI agent and therefore mainly share the same aims. Finally, if it's actually true that discrepancies in ethical aims mostly reflect discrepancies in beliefs about non-moral claims, and if AI agents would typically be more rational than people, then AI agents might also tend to converge more readily than people do.

B.2.6.3 The Overall Effect of Changes in Culture and Values

Ultimately, it does seem like we should expect culture and values to continue to change. It also seems like, in a world with radical AI capabilities, these changes would at least in part be downstream of AI progress. But I have close to no idea what the changes will be. I also do not see any strong reason to *trust* that they will tend to be pacifying.

I think that the strongest argument for expecting a pacifying effect is that there could be a tendency toward *convergence* in actors' ethical aims. Convergence would then seem to limit the potential for conflict. However, in my view, the arguments in favor of strong

convergence are still fairly tenuous.

B.3 Conclusion

This final section of the thesis has been especially speculative. (I imagine that, after reading the previous thirty pages, the disclaimer now goes without saying.) I have tried to look forward into a murky possible future, one that contains truly radical advances in AI technology, and then engage in causal reasoning about several specific mechanisms that could lead the rate of conflict to change.

Overall, while there were considerations that pushed in both directions, my sense has been that there are overall *stronger* reasons to expect AI to have a pacifying effect. Although I will not go through the exercise of updating our overall conflict forecast again (leaving it at a 15% probability of a peaceful future), I now suspect that reasoning about the various possible effects of AI actually *should* lead us to become meaningfully more optimistic. Our sense of the "spectrum of reasonable disagreement" should expand, but our central estimate should also — I would guess — increase.

My feeling is also that there has been intellectual value in this extended exercise, beyond any usefulness it has for forecasting the future of conflict. By asking how the risk of conflict could change, if various fundamental and usually taken-for-granted parameters change, we have also been forced to develop a fuller understanding of key theories regarding the causes of war.