Economic Quarterly—Volume 105, Number 2—Second Quarter 2019—Pages 59–103

Time-Varying Skewness and Real Business Cycles

Lance Kent and Toan Phan

A growing literature in macroeconomics and finance has found important economic effects of variations in risk, in particular shocks to the *volatility* of key macroeconomic variables (such as total factor productivity). However, much less is known about the importance of shocks to the *skewness* of macroeconomic variables.[1]

In this paper, we seek to quantify the economic effects of skewness shocks. To this end, we augment a small open economy real business cycle model with a novel feature: discrete regime changes in the higher-order moments of exogenous shocks, modeled as shocks to total factor productivity (TFP). We assume that in each period the economy can be in one of two possible Markov states: an *unrest* state or a *quiet* state. The unrest state is assumed to be associated with a substantial increase in volatility and negative skewness of shocks. This assumption is motivated by our empirical findings about the moments of business cycles of many countries that experience political unrest (see the discussion of our calibration below). Hence, unrest is effectively a shock to the second-order and third-order moments of the distribution of economic shocks.

Amazon, lanckent@amazon.com (work completed while the author was affiliated with the College of William & Mary); Federal Reserve Bank of Richmond, toanvphan@gmail.com. This paper supersedes an earlier working paper entitled "Business Cycles and Revolutions." The authors are grateful for helpful discussions with Daron Acemoglu, Filipe Campante, Matthias Doepke, B.R. Gabriel, Elias Papaioannou, and Romain Wacziarg. The authors would also like to thank seminar participants at the Atlanta Fed, Midwest Macro, the National University of Singapore, the 2014 ASSA Meetings, Duke University, the 2014 ISNIE meetings, the 2014 ESEM meetings in Toulouse, Queen's University Belfast, Utrecht University, and Maastricht University for their suggestions and feedback. The views expressed herein are those of the authors and not those of the Federal Reserve Bank of Richmond or the Federal Reserve System.

[1] See the related literature section for a discussion of existing studies.

To solve the model, we develop a third-order perturbation method to approximate the endogenous reactions to shocks to the second-order and third-order moments of TFP. Existing methods to solve and simulate models (including global approximations to policy functions as in Judd [1996] or Richter et al. [2014] or perturbation methods as in Andreasen et al. [2017]) rely on Monte Carlo simulations to calculate the dynamics of third-order moments of endogenous quantities such as output and consumption. However, Monte Carlo simulations are problematic for the computation of higher-order moments such as skewness because these higher-order moments are more sensitive to simulation error.[2] To overcome this problem, we build upon the method of Andreasen (2017) to calculate generalized impulse response functions (GIRF) of third-order approximations of third-order moments of endogenous variables. Our solution method exploits computational symbolic algebraic manipulation to calculate the third-order moments *without* Monte Carlo simulations. This technical innovation is nontrivial, since it requires solving for the dynamics of over 20,000 polynomials, in the presence of a Markov-switching state, that are up to ninth order in the state variables. Furthermore, our approach is readily applicable to other DSGE models, especially those for which the dynamics of higher-order moments of endogenous variables are of interest.

Calibration: To calibrate the model, we document and exploit the substantial changes in higher-order moments of aggregate economic variables during periods of mass political unrest. Unrest episodes, which are well-documented by the political science literature (Chenoweth and Lewis 2013), are helpful in identifying higher-order moment shocks for several reasons. First, we find that these episodes are associated with substantial increases in the volatility and negative skewness of growth rates of output, consumption, and investment. For instance, on average, a year during an unrest episode is associated with a more than 50 percent increase in the volatility and a more than three times increase in the negative skewness of output growth. The changes in higher-order moments of aggregate variables (output growth, consumption growth, and investment growth) associated with an episode of unrest can be estimated with reasonable precision, since the database provides a relatively large number of country-year observations (with

[2] The calculation of skewness and other higher-order moments is sensitive to the tails of the distribution of interest. Since realizations on the tails are rare, many Monte Carlo draws are needed to ensure that the tails are sufficiently sampled. Therefore, for a given simulation length, the influence of Monte Carlo simulation error is going to be much more pernicious for higher-order moments, such as skewness, than for lower-order moments, such as the mean.

eighty-four unrest episodes between 1960 and 2006, each lasting more than five years on average).

Second, since the model assumes that shocks are common knowledge, we ideally want to identify shocks using events that are easily observed for all agents, at home or abroad. Mass unrest episodes are appropriate for this end, as they are major events, and agents in the economy as well as investors abroad do not need to be econometricians to learn that a campaign of mass political unrest is underway. Hence, the onset of an unrest episode is likely to have a direct effect on economic agents' perceptions of risk. Furthermore, since the impulse response exercises assume unanticipated shocks, we ideally want to use events that are ex-ante difficult to predict. Unrest episodes are again appropriate for this end, as it has been well-documented that mass unrest is largely unanticipated because it requires unpredictable shocks that enable a large number of nonstate actors to overcome informational and coordination problems.[3]

Results: Our model shows that the increase in volatility and especially negative skewness when the economy enters an episode of unrest has quantitatively substantial impacts on economic activities. In the baseline calibration, the observed changes in volatility and negative skewness can explain 21 percent of the observed drop in average output growth, 45 percent of the drop in average consumption growth, and 51 percent of the drop in average investment growth during unrest episodes. More importantly, *the increase in negative skewness accounts for about half of these drops in growth.*

Intuitively, when shocks become more negatively skewed, risk-averse agents know that realizations on the left tail of the distribution of shocks have become more likely. The increase causes agents to shift their portfolios to safer assets abroad and accumulate stocks of these safer assets, leading to capital outflow and drops in domestic investment and output. The consequences of this increased mass on the left tail are heightened under Epstein-Zin preferences. A Taylor expansion of the household's Bellman equation reveals that Epstein-Zin preferences punish and reward, respectively, the second and third central moments of the future value function. To a second-order approximation, Epstein-Zin preferences penalize the second central moment, i.e., variance. To a third-order approximation, the preferences gain an additional term that rewards the third central moment, which is the product of skewness and variance raised to the power of 3/2. Therefore, the quantitative

[3] See, e.g., Kuran (1989), Chenoweth and Stephan (2011), and Edmond (2013). We also verify this in our probit analysis to predict the onset of unrest in Appendix A.2.

effects of time-varying variance is amplified by time-varying negative skewness.

We demonstrate the quantitative significance of skewness by comparing the losses in economic activities during unrest under a second-order approximation to the same losses under a third-order approximation. The second-order approximation can account for only half of the economic losses that the third-order approximation can. Therefore, negative skewness is revealed as an important component of risk.

Related literature. Our paper is related to several strands of the literature on higher-order moments of business cycles. First, there is a growing body of research that emphasizes the importance of the time-varying volatilities of economic variables (e.g., Justiniano and Primiceri 2008; Caldara et al. 2012; Arellano et al. 2012; Christiano et al. 2014; and Gilchrist et al. 2014. The study that is the closest to ours in quantifying the impact of time-varying higher-order moments is Fernández-Villaverde et al. (2011). They consider a stochastic volatility process for the real interest rate and explore the impacts of interest rate volatility shocks to economic activities. The primary difference between our paper and this literature is that while they focus only on shocks to second moments, we focus on shocks to both second-order and third-order moments.

Second, there is a related body of macro-finance research that stresses the importance of skewness (e.g., Rancière et al. 2008; Barberis and Huang 2008; Guvenen et al. 2014; Salgado et al. 2015; Feunou et al. 2015; and Colacito et al. 2015). Our analysis is most related and complementary to that of Colacito et al. (2015), who show the importance of time-varying skewness in a macro-finance model with Epstein-Zin preferences. The major difference is that while they focus on the effects of skewness on financial variables (implied equity Sharpe ratios and equity risk premia), we focus on the effects on real economic variables (the growth rates of output, consumption, and investment). Also, while they focus on the United States, we focus on emerging and developing economies. Finally, while they calibrate the model by looking at analysts' forecasts for the U.S. economy, we look at the changes in higher-order moments of real economic variables during unrest episodes.

Finally, our paper is also related to a body of literature that emphasizes the importance of rare disasters in explaining macroeconomic phenomena (e.g., Barro 2006; Gourio 2012; Andreasen 2012; Gabaix 2012). A key insight from this literature is that variations in the probability of rare disasters, modeled as events on the far left tail of the distribution of shocks, can have first-order macroeconomic effects, as they influence the precautionary behaviors of risk-averse agents. Our

paper points out that time-varying negative skewness has similar effects. This is because an increase in negative skewness implies a higher probability of states with very low consumption. However, our estimation approach is different and complementary to existing approaches in this literature. Since rare disasters occur infrequently in data, the literature usually does not estimate the time variation in the probability of disasters from data,[4] or it employs calibrations to proxies such as time-varying volatility of equity returns (e.g., Gourio et al. 2013). In contrast, we exploit the uncertainty associated with episodes of unrest to estimate the time variation in the skewness of economic shocks when the economies enter and exit unrest.[5]

Our paper is organized as follows. Section 1 describes our data sources and documents several stylized facts on business cycles during unrest episodes. Section 2 introduces unrest to a standard small open economy model and calculates how much of the stylized facts can be explained by changes in the distribution of shocks. Section 3 concludes.

1. DATA AND STYLIZED FACTS

Data Sources and Definitions

For economics and other data, we use annual panel macroeconomic data from 154 countries listed in the World Bank's World Development Indicators (WDI) database over the interval 1960-2006. This includes three time series: real output, real investment, and real consumption. We also use WDI data on the Gini coefficient and Alesina et al.'s (2003) data on ethnic, linguistic, and religious fractionalization as control variables

For mass unrest episodes, we use the Nonviolent and Violent Campaigns and Outcomes (NAVCO) dataset, version 2.0 (Chenoweth and Lewis 2013). NAVCO 2.0 provides a "consensus population" of all known continuous and large (having at least 1,000 observed participants) organized unrest campaigns between 1945 and 2006[6] that satisfy a series of conditions, as detailed in Appendix C. Each episode has an onset year and an end year. The onset year is defined as the first year with a series of coordinated, contentious collective actions with at least 1,000 observed participants. The episode is recorded as over when peak

[4] E.g., Nakamura et al. (2013) allow disasters to be correlated across countries but suppose that the probability of a given country entering into a disaster "on its own" is fixed over time.

[5] It is important to note that we do not identify unrest episodes themselves as disasters.

[6] More recent versions of the dataset include more recent years.

Figure 1 Example of Business Cycles Around Unrest

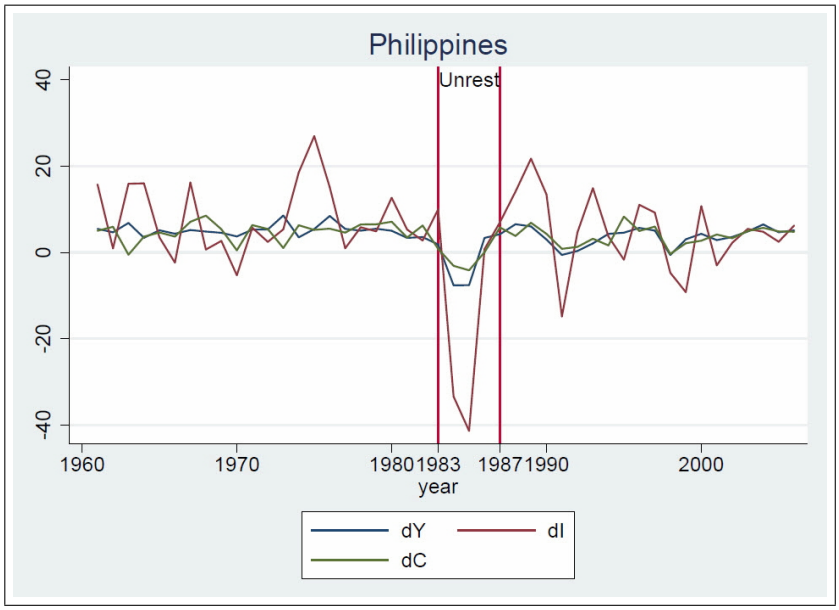

Notes: Growth rates of output, consumption, and investment (dY, dC, and dI respectively) of the Philippines around the People's Power Revolution (1983-87).

participation drops below 1,000. Overall, the NAVCO dataset provides 157 episodes of nonviolent and violent mass political unrest around the world between 1945 and 2006. Of these, there are eighty-four episodes in the years between 1960 and 2006, the period for which we have both unrest and economic data. Over this period, the average duration of an episode is 5.99 years.

Examples include many pro-democracy movements of civil unrest in Latin America, the Philippines's People Power Revolution (1983-87), Indonesia's civil unrest against Suharto (1997-98), and Mozambique's RENAMO resistance movement (1979-1992); for a complete listing of these episodes, see Appendix A.1. As an illustration, Figure 1 plots the time series of the growth rate in aggregate economic variables for the Philippines around the People's Power Revolution.

Stylized Facts on Business Cycles During Unrest

We now investigate the relationship between unrest and macroeconomic activities. The goal of this section is to arrive at a set of moments that will be used as calibration targets for the structural model of the following section. We focus on the contemporaneous association between unrest in a given country-year and the growth rates of output, consumption, and investment. We follow others in the macroeconomic literature (e.g., Fernández-Villaverde et al. 2011) and do not explicitly model why the higher-order moments change, nor do we attempt to make any causal claims about the contemporaneous causal impacts of unrest on output or vice versa.

We calculate the growth rates of output, consumption, and investment by the first difference in logs of the variable at constant 2005 USD and then remove a country-specific average growth rate from each series. That is, if the real output for country i in year t is Y_{it}, then we calculate the raw growth rate as $\Delta Y_{it} \equiv 100(\ln Y_{it} - \ln Y_{it-1})$. Then we take out the country's average growth rate to yield a demeaned output growth rate of $g_{Y,it} \equiv \Delta Y_{it} - \frac{1}{T_i}\sum_t \Delta Y_i$. A similar method is applied to demean consumption and investment growth. We demean to isolate fluctuations at the business cycle frequency and to control for differences in country-specific average growth rates.

We then contrast the distributions of growth rates during unrest ($g_{it}|U_{it} = 1$) against moments during quiet times of no unrest ($g_{it}|U_{it} = 0$) in Figure 2. The left column of Figure 2 displays smoothed kernel estimates of the empirical probability density functions for the growth rates of output, consumption, and investment, and the right column displays the corresponding empirical cumulative distribution functions. The probability density functions are estimated by Epanechnikov kernels with a bandwidth of 2 percentage points for output and consumption, and 4 percentage points for investment. The figures suggest that the distributions of the growth rates are more negatively skewed during unrest episodes.

To have numerical comparisons, Table 1 displays the means, standard deviations, skewnesses, and kurtoses of (country-demeaned) output growth, consumption growth, and investment growth during and outside of unrest episodes. All confidence intervals are bootstrapped with 500 replications and are reported at the 95 percent level. The first two columns report the estimated moments. The third column reports the difference in the estimated moments, along with the p-value for a test of the null hypothesis that there is no difference between the corresponding moments. The fourth column reports the ratio of the estimated standard deviations, along with the p-value for the Levene

Figure 2 Distribution Functions of Macro Variables

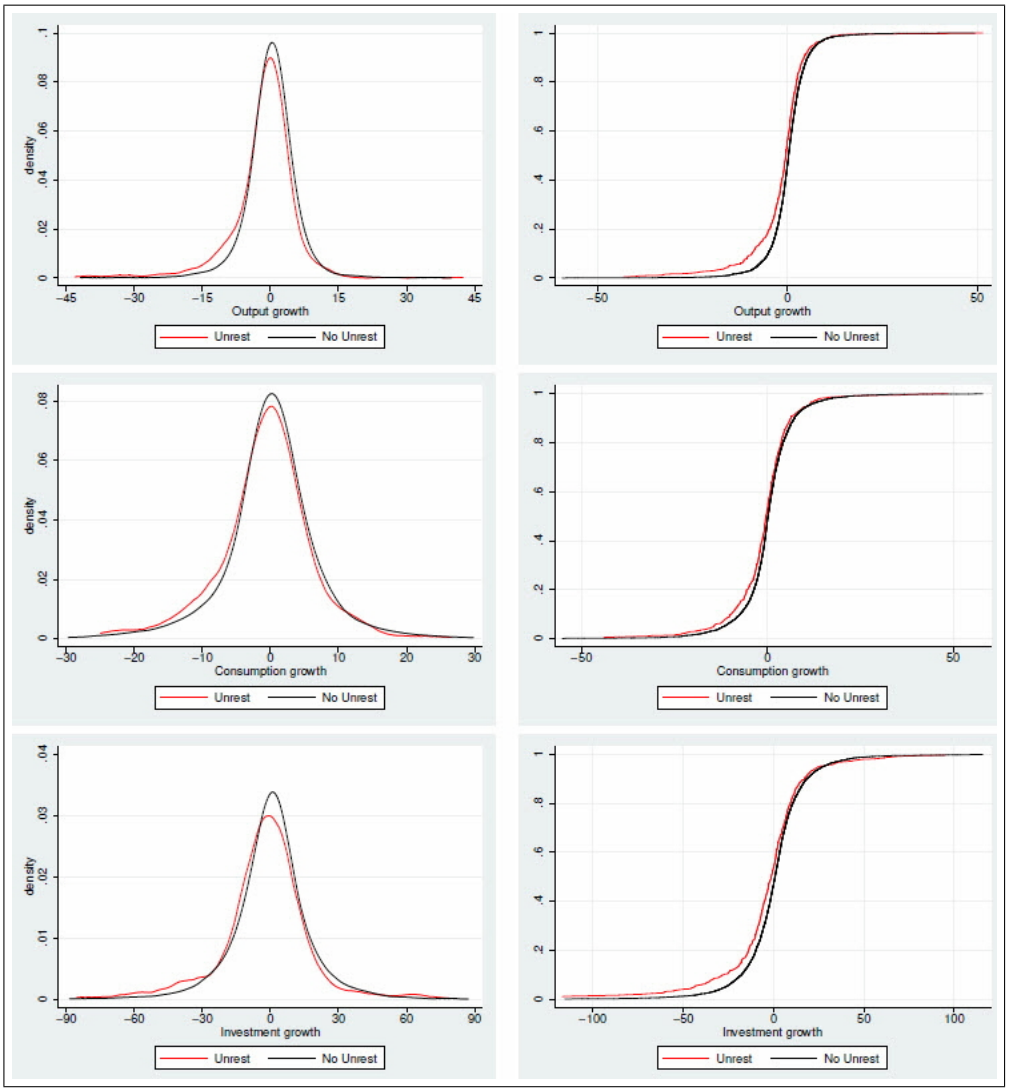

Notes: Smoothed empirical probability density functions and empirical cumulative distribution functions of output, consumption, and investment growth, unrest vs. no unrest.

test of the equality of variances. The fifth column reports the p-value for the Kolmogorov-Smirnov test of whether the two distributions of shocks (under unrest and no unrest) are the same.

Table 1 Empirical Moments In and Out of Unrest

	No unrest (c.i.)	Unrest (c.i.)	Difference [p-value]	Ratio [p-value]	K-S test [p-value]
Output growth					[0.00]
Mean	0.17	-1.75	-1.92		
	(0.03,0.32)	(-2.46,-1.04)	[0.00]		
Standard Dev.	5.62	8.63		1.53	
	(5.30,5.95)	(7.05,10.20)		[0.00]	
Skewness	-0.69	-2.26	-1.57		
	(-1.68,0.30)	(-4.16,-0.35)	[0.15]		
Kurtosis	23.61	23.38	-0.23		
	(15.95,31.27)	(13.86,32.91)	[0.97]		
Consumption grth					[0.00]
Mean	0.12	-1.10	-1.22		
	(-0.11,0.36)	(-1.90,-0.31)	[0.00]		
Standard Dev.	8.34	9.80		1.17	
	(7.69,8.99)	(7.30,12.29)		[0.00]	
Skewness	0.40	-2.96	-3.36		
	(-1.38,2.17)	(-6.52,0.60)	[0.10]		
Kurtosis	33.46	39.88	6.42		
	(17.34,49.57)	(6.66,73.10)	[0.73]		
Investment grth					[0.00]
Mean	0.40	-3.56	-3.96		
	(-0.19,1.00)	(-6.00,-1.12)	[0.00]		
Standard Dev.	20.29	27.45		1.35	
	(18.69,21.89)	(22.93,31.96)		[0.00]	
Skewness	-0.94	-0.71	0.24		
	(-2.52,0.64)	(-2.14,0.73)	[0.83]		
Kurtosis	31.78	15.11	-16.66		
	(17.77,45.80)	(11.21,19.00)	[0.03]		

Notes: Empirical moments in and out of unrest with bootstrapped 95 percent confidence intervals (in brackets) and p-values (in square brackets) on hypothesis tests that there is no difference between the two distributions. The first two columns report the estimated moments. The third column reports the difference in the estimated moments, along with the p-value for a test of the null hypothesis that there is no difference between the corresponding moments. The fourth column reports the ratio of the estimated standard deviations, along with the p-value for the Levene test of the equality of variances. The fifth column reports the p-value for the Kolmogorov-Smirnov test of whether the two distributions of shocks (under unrest and no unrest) are the same.

Table 1 shows that a period of unrest is associated with significant losses in growth. The per-year loss in output growth (relative to periods without unrest) is 1.92 percent, statistically significant at the 1 percent level. This estimated per-year loss is nontrivial, especially given that unrest is persistent once started. The estimated cumulative loss is relatively substantial at 11.50 percent of the base (pre-onset)

year's output.[7] The annual loss in consumption growth is 1.22 percent, which is smaller than that of output growth. At the same time, investment growth losses are larger than output growth, at 3.96 percent. In cumulative terms, consumption and investment losses amount to 7.31 percent and 23.71 percent, respectively. Note that this ordering of the loss in investment, output, and consumption is consistent with the permanent income hypothesis, which predicts that investment is more sensitive to shocks than output, which is in turn more sensitive than consumption.

Furthermore, Table 1 shows that the standard deviations of the growth rates of output, consumption, and investment substantially increase during unrest episodes. The fourth column of Table 1 displays the ratio of standard deviations. We can see that the standard deviation of output growth is 53 percent larger in unrest, and the standard deviations of consumption and investment growth are 17 percent and 35 percent larger, respectively. The column also reports the p-values of Levene's test of equality of variances between various forms of unrest against the baseline of no unrest. The p-values show that all of these increases are highly statistically significant: well below 0.01 for all three.

Table 1 also shows that both output and consumption growth becomes more negatively skewed during unrest. The difference in the skewness between unrest and no unrest is -1.57 for output growth and -3.36 for consumption growth. The bootstrapped p-value for the hypothesis that the difference in skewness is equal to zero is 0.15 for output growth and 0.10 for consumption growth. While it is generally difficult to estimate higher-order moments of relatively infrequent events with great confidence, we believe that these differences in skewness are economically significant. The greater variance and larger left tail of many distributions are also visually discernible in Figure 2.[8] This discernible mass on the left tail corresponds to a continuous range from moderately to extremely bad outcomes. The difference between a period of unrest and a period with no unrest then is not the increased probability of a single disaster but an increase in the probability of a whole range of bad outcomes.

[7] If p is the continuation probability and x is the annual loss, then the cumulative loss is estimated to be $\frac{x}{1-p}$.

[8] While there is also a visibly larger left tail for the distribution of investment growth, the bootstrapped difference in the skewness in investment growth between unrest and no unrest is not significantly different from zero, with a p-value of 0.83. This is because there are a few observations of investment growth that are very large in absolute value on both sides of the distribution (consistent with sharp falls in investment and subsequent rebounds), and the bootstrapped estimate of the difference in skewness is sensitive to these outliers.

Finally, as the fifth column of Table 1 shows, under the Kolmogorov-Smirnov test, we can reject the hypothesis that the two distributions of shocks (under unrest and under no unrest) are the same, as the associated p-value is zero for each series (output growth, consumption growth, or investment growth).

We summarize our results in the following stylized fact:

Fact: *Episodes of mass political unrest are associated with statistically and economically significant economic costs: the distributions of output, investment, and consumption growth during unrest have lower means and higher variances than the distributions in periods of no unrest. In addition, the distributions of output and consumption growth are more negatively skewed during unrest.*

One potential mechanism that could explain the increased volatility and negative skewness in economic activities is that unrest is associated with substantial increases in the probability of institutional disruptions. In Appendix A.3, we document that the probabilities of large political and government changes, including major changes in polity and coups, substantially increase during unrest episodes. Large political changes are often associated with significant changes in legal and economic institutions, such as the protection of property and investment, which are key determinants of investment and growth (Acemoglu and Robinson 2005; and Acemoglu et al. 2014). Therefore, unrest episodes can increase the probability and severity of economic disasters.

2. QUANTITATIVE ANALYSIS

Model

How much of observed declines in average output, consumption, and investment growth during unrest, as reported in the previous section, can be attributed to volatility and skewness shocks? To answer this question, we augment a standard small open economy with a regime-switching process for the volatility and skewness of TFP. We calibrate the regime-switching process to moments that were estimated from data in the previous section.

Consider a canonical small open economy model with a representative household. Domestic firms competitively produce a numeraire good Y_t using capital K_{t-1} and labor H_t, subject to TFP ζ_t:

$$Y_t = \zeta_t K_{t-1}^{\alpha} (H_t)^{1-\alpha}.$$

These firms take factor prices R_t and W_t as given. Their first-order conditions on their optimal choices of capital and labor equate these

factor prices with the corresponding marginal products in production:

$$\begin{aligned} W_t &= (1-\alpha)\zeta_t K_{t-1}^{\alpha} H_t^{-\alpha} \\ R_t &= \alpha\zeta_t K_{t-1}^{\alpha-1} H_t^{1-\alpha}. \end{aligned}$$

Unrest shock. We introduce a regime-switching process. Let u_t be an exogenous two-state Markov process, with $u_t = 1$ representing the country being in unrest in period t and $u_t = 0$ representing no unrest, or a quiet time, in period t. Transitional probabilities are calibrated to match the probability of unrest onset and the persistence of unrest observed in data.

To model how unrest affects economic activities in the most tractable way, we assume that unrest affects the TFP process. Intuitively, as unrest episodes are associated with significant economic and political instability, they will affect the productivity of many economic sectors by, for instance, affecting the efficiency of resource allocation (Acemoglu et al. 2014). Such effects can be captured in a reduced form by a wedge to TFP, as in Chari et al. (2007).

Remark. Recall that our goal is to analyze the extent to which the shocks to higher-order moments of aggregate macroeconomic variables that we observe during unrest can explain the observed average losses in output, consumption, and investment growth. To conduct this analysis in the simplest and clearest possible way, we assume that unrest is a shock *only* to higher-order moments of the TFP process and not to the first moment. Obviously, this is a simplifying assumption and will likely lead to underestimations of the economic impacts of unrest. The model can be extended to allow for the possibility that unrest affects the first moment as well, but this will complicate the analysis. We will show that, even without an immediate associated fall in average productivity, a higher-order moment shock is enough to generate large changes in macroeconomic aggregates in line with the data.

Specifically, assume that TFP ζ_t consists of a growth component $(g^t)^{1-\alpha}$ and a level component A_t:

$$\zeta_t = (g^t)^{1-\alpha} A_t,$$

where, for numerical simplicity, we have assumed that growth rate is a constant g. However, level component A_t follows an autoregressive process with autoregressive parameter ρ and i.i.d. shocks ε_t:

$$\ln A_t = \rho \ln A_{t-1} + \varepsilon_t.$$

The stochastic process for ε_t depends on whether the economy is currently experiencing unrest. While in unrest ($u_t = 1$), shock ε_t is distributed Normal Inverse Gaussian with mean 0, standard deviation σ_u, skewness s_u, and kurtosis κ_u. While not in unrest ($u_t = 0$), shock ε_t is distributed Normal Inverse Gaussian with mean 0, standard devia-

tion σ_q, skewness s_q, and kurtosis κ_q. The Normal Inverse Gaussian distribution has been used in the finance literature to model skewed distributions with fat tails (e.g. Barndorff-Nielsen 1997; Andersson 2001; and Mencía and Sentana 2012). The fact that the mean of ε_t is the same whether $u_t = 0$ or $u_t = 1$ reflects the assumption that unrest only affects higher-order moments of TFP.[9]

Preferences: As is now standard in the macro-finance literature (e.g., Gourio 2012; and Colacito and Croce 2013), we assume the representative household has recursive preferences as in Epstein and Zin (1989). These preferences allow us to distinguish between the intertemporal elasticity of substitution and risk aversion (captured by ζ and γ below). Moreover, these preferences nest the standard expected utility with constant relative risk aversion (CRRA) as a special case.

Let C_t denote household consumption in period t, and let $\tilde{C}_t \equiv C_t - \theta\omega^{-1}g^{t-1}H_t^{\omega}$ denote labor-adjusted consumption, where θ and ω are preference parameters. Then, we follow the sign convention of Rudebusch and Swanson (2012) and define the representative household's preferences as:

$$\tilde{V}_t \equiv \begin{cases} \left((1-\beta)\tilde{C}_t^{1-\varsigma} + \beta E_t\left[V_{t+1}^{1-\gamma}\right]^{\frac{1-\varsigma}{1-\gamma}}\right)^{\frac{1}{1-\varsigma}} & \text{if } \tilde{C}_t^{1-\varsigma} \text{ is always positive} \\ \left((1-\beta)\tilde{C}_t^{1-\varsigma} - \beta E_t\left[(-V_{t+1})^{1-\gamma}\right]^{\frac{1-\varsigma}{1-\gamma}}\right)^{\frac{1}{1-\varsigma}} & \text{if } \tilde{C}_t^{1-\varsigma} \text{ is always negative} \end{cases} \tag{1}$$

This convention ensures that the value function and the instantaneous payoff have the same sign.

Households supply capital and labor to the domestic firms, consume domestic goods, invest subject to an adjustment cost in capital, and trade noncontingent bonds in the international credit market:

$$V_t = \max_{C_t, D_t, I_t, H_t} \tilde{V}_t,$$

subject to:

$$\begin{aligned} C_t + D_{t-1} + I_t + \frac{\phi}{2}\left(\frac{K_t}{K_{t-1}} - g_t\right)^2 K_{t-1} &= R_t K_{t-1} + W_t H_t + \frac{D_t}{1+r_t} \\ K_t &= (1-\delta)K_{t-1} + I_t. \end{aligned}$$

[9] The unrest shock in each period t affects the distribution of the TFP in period t, and because TFP is autocorrelated, the unrest shock will affect the distribution of future TFP terms too. This is different from a "news shock" that does not affect current TFP, only future TFP.

We assume that the interest rate households borrow at is a function of the aggregate stock of debt D_t:

$$r_t = r^* + \psi(e^{D_t/g^t - \bar{d}} - 1),$$

where r_t is the interest rate, r^* is a constant representing the world's risk-free interest rate, and $\bar{d}$ and ψ are exogenous constants. This debt-elastic interest rate is a standard assumption to ensure that the equilibrium is stationary (e.g., Schmitt-Grohé and Uribe 2003).

Finally, a recursive equilibrium is defined as a set of policy functions for C_t, V_t, K_t, D_t, Y_t, r_t, I_t, H_t, W_t, and R_t as functions of K_{t-1}, D_{t-1}, A_t, and u_t such that all agent expectations are rational and the optimality conditions, constraints, and laws of motion described above hold.

Solution Method

One way to derive moments of output, consumption, and investment growth from the model is to simulate a very long time series in which the country transitions into and out of unrest with the same probabilities as in the data. But since unrest is rare, we would need an extraordinarily long simulated time series to reduce the Monte Carlo noise around our estimates of those higher-order moments. Instead, we adapt the pruning method from Andreasen et al. (2017) to get closed-form solutions for the paths of conditional moments of endogenous variables, the GIRF. We first describe how we calculate a GIRF and then how we use the GIRF to compare the model against the data. All details on the computational strategy, from approximation to pruning and the GIRF, are given in the Appendix.

We define the GIRF as follows. Let y_t denote the log-deviation of output Y_t from its steady-state value. Then Δy_t is the growth rate of output Y_t. Let X_t denote a vector of the first three powers of the growth rates of output, consumption, and investment:

$$X_t \equiv \left(\Delta y_t, \Delta i_t, \Delta c_t,\ (\Delta y_t)^2, (\Delta i_t)^2, (\Delta c_t)^2,\ (\Delta y_t)^3, (\Delta i_t)^3, (\Delta c_t)^3\right).$$

The GIRF is the evolution over time of the difference of conditional expectations of X_t between two conditioning sets, differing with respect to two given time series of realizations of unrest, $u = \{u_t, -\infty < t < \infty\}$ and $\tilde{u} = \{\tilde{u}_t, -\infty < t < \infty\}$:

$$GIRF^u(X_t) \equiv E[X_t|u] - E[X_t|\tilde{u}].$$

The first path, u, represents a country that starts with no unrest and then enters into unrest at $t = 1$ and stays there. That is, $u_t = 0$

$\forall t \leq 0$ and $u_t = 1$ $\forall t \geq 1$. The second counterfactual path, $\tilde{u}$, is one where the country never enters unrest: $\tilde{u}_t = 0$ $\forall t$.

Remark. The GIRF is useful for our purposes for several reasons. First, we want to calculate the moments that would be uncovered from a simulation. The conditional expectations in the GIRF allow us to consider the effects of shocks over the course of the GIRF. This is important, since under a nonlinear approximation to the policy function, the presence of shocks will cause the ergodic moments of all variables to differ from those in the absence of shocks. Second, since X_t contains powers and products of endogenous variables, we can find paths not just for conditional means, but also for conditional variances and skewnesses of the endogenous variables of interest given the paths for the components of X_t. Moreover, the GIRF allows us to avoid measurement error, which is a problem for estimating higher-order moments of simulated series from a finite simulation length. While Andreasen et al. (2017) rely on SMM for higher-order moments, we use the computer algebra software Mathematica to calculate GIRFs for these moments symbolically, term by term, and avoid Monte Carlo error.

The GIRF provides the conditional moments in the first year of an unrest episode, the second year, and so on. The moments from the data presented in the previous section are *weighted* averages over the years in observed unrest episodes because years that are closer to the beginning of an episode are more likely observed than years that are many years after the beginning of an episode. If $p = \Pr(U_t|U_{t-1} = 1)$, then the probability of a given observed year of unrest being the nth year of unrest ($n \geq 1$) within its respective episode is $(1-p)p^{n-1}$. Thus, to construct the single value for average value of X on unrest, we take a weighted average of a GIRF where a country enters into unrest and stays there but with smaller and smaller weight given to later periods of unrest. That is, we calculate $\sum_{t=1}^{\infty}(1-p)p^{t-1}GIRF^u(X_t)$.

Calibrations

First, we calibrate the model's basic parameters using standard values from the small open economy literature. These numbers are listed in the top panel of Table 2.[10] We allow the values for Epstein-Zin

[10] The sensitivity parameter of interest rate to debt is simply set to a small value to avoid a unit root, as in García-Cicco et al. (2010) and Schmitt-Grohé and Uribe (2003).

Table 2 Calibrated Parameters

	Parameter	Value	Source/Target
	From literature		
α	Capital share in production	0.32	García-Cicco et al. (2010)
β	Discount factor	0.922	–
δ	Depreciation	0.126	–
ϕ	Adjustment costs to capital	3.3	–
g	Trend growth rate	1.005	–
θ	Disutility from labor	0.224	–
ω	Disutility from labor	1.6	–
$\bar{d}$	Steady-state debt level	0.007	–
ψ	Interest rate sensitivity to debt	10^{-5}	–
ς	Inverse intertemporal elasticity of substitution	0.9 to 5	Table 3
γ	Risk aversion	5 to 20	Table 3
	Estimates from data		
p_{onset}	Probability of unrest onset	0.014	Appendix A.2
$p_{cont.}$	Probability of unrest continuation	0.833	Appendix A.2
	Chosen to match target		
σ_q	Std. of TFP shock ε_t in quiet times	2.75	Table 1
s_q	Skewness of TFP shock ε_t in quiet times	-1.10	Table 1
κ_q	Kurtosis of TFP shock ε_t in quiet times	22	(*)
σ_u	Std. of TFP shock ε_t during unrest	4.66	Table 1
s_u	Skewness of TFP shock ε_t during unrest	-2.83	Table 1
κ_u	Kurtosis of TFP shock ε_t during unrest	22	(*)

Notes: (*) means chosen sufficiently high to permit existence of Normal Inverse Gaussian distribution.

Table 3 Epstein-Zin Parameter Calibrations in the Literature

	ς	γ
Fernández-Villaverde et al. (2011)	5	5
Colacito et al. (2013)	0.9	10
Vissing-Jørgensen and Attanasio (2003)	0.9	20

preference parameters to vary within the standard range of values of the literature, surveyed in Table 3.[11]

[11] In Vissing-Jørgensen and Attanasio (2003), the estimated risk-aversion parameter can take a wide range of values, as large as thirty. To be conservative, we only set the maximum risk-aversion to be twenty.

Second, we calibrate parameters for the unrest process and the higher-order moments of TFP innovation ε_t to estimated moments from our empirical analysis in Section 1. Care must be paid to the calibration of the higher-order moments of TFP, both in unrest and in quiet times. The parameters chosen in the model govern the exogenous TFP process, but they are chosen to match the moments of endogenous quantities. It is relatively straightforward (as one could even rely on closed-form solutions) to choose the volatility of a shock process given a desired volatility of an endogenous quantity, such as output growth under a log-linear approximation to equilibrium. However, it is much less straightforward to choose higher-order moments of a shock process to match higher-order moments of a nonlinear approximation of the law of motion for an endogenous variable. Therefore, the parameters σ_q, σ_u, s_q, and s_u are chosen so that the ergodic standard deviation and skewness of output growth, and the average generalized impulse responses of the standard deviation and skewness of output growth, match those in the data.[12]

Results

Model's performance relative to data. We compare the average loss in output, investment, and consumption growth from the GIRF $\sum_{t=1}^{\infty}(1-p)p^{t-1}GIRF^u(\Delta y_t)$ to the corresponding observed average loss in growth as documented in Section 1. Table 4 reports the percentage of observed growth loss that can be explained by the calibrated model. The overall effect is an endogenous response of endogenous variables to an unrest shock that increases the volatility and negative skewness of TFP shocks, with an interplay of capital adjustment costs and preferences over the time resolution of risk. In each panel, we report the percentage obtained by using the first-, second-, and third-order approximations of the solution to the model. Note that by construction, the percentage explained using a first-order approximation is zero, as we assume that unrest does not affect the first moment of TFP shocks. The columns report the results with different preference parameters.

Table 4 shows that, under the baseline specification (the first column), the model explains 21 percent of the average output growth loss, 45 percent of the average consumption growth loss, and 51 percent of the average investment growth loss. This amounts to an output growth

[12] We do not attempt to match kurtoses exactly, since we approximate equilibrium only to third order. We choose the kurtosis of the TFP processes high enough to permit existence of the Normal Inverse Gaussian distribution for the calibrated second-order and third-order moments. The calibrated kurtosis of TFP is approximately equal to that of the empirical distribution of output growth.

Table 4 Numerical Results

	Numerical Results		
	Baseline $\varsigma = 0.9, \gamma = 10$	High risk av. $\varsigma = 0.9, \gamma = 20$	No EZ, low risk av. $\varsigma = 5, \gamma = 5$
Output growth			
First order	0	0	0
Second order	11	21	6
Third order	21	62	9
Consumption growth			
First order	0	0	0
Second order	22	742	7
Third order	45	128	11
Investment growth			
First order	0	0	0
Second order	27	51	16
Third order	51	148	23

Notes: Numerical results for the percentages of the empirically observed average losses in the growth rates of output, consumption, and investment that are explained by the model. The rows show the percentage explained by using first-order, second-order, and third-order approximations of the solution to the model.

loss of 0.40 percent per year, a consumption growth loss of 0.55 percent per year, and an investment growth loss of 2.01 percent per year. In cumulative terms over the average episode duration, this is an output growth loss of 2.41 percent, a consumption growth loss of 3.28 percent, and an investment growth loss of 12.09 percent.

The second column of Table 4 shows that, not surprisingly, the model can explain more with a larger coefficient of risk aversion ($\gamma = 20$ instead of $\gamma = 10$). There, the fractions of growth losses explained increase to 62 percent for output, 128 percent for consumption and 148 percent for investment (thus this calibration "overexplains" the losses in consumption and investment). On the other hand, when we shut down Epstein-Zin preferences and use a lower coefficient of risk aversion (the third column), the fractions of growth losses explained decrease to 9 percent, 11 percent, and 23 percent for output, consumption, and investment, respectively.

It is not surprising that the model cannot fully explain the observed losses, since we assume that unrest only affects higher-order moments of TFP shocks, and not the first-order moment, thus abstracting away

from factors such as reallocation of resources between sectors of the economy that may directly affect the average productivity.[13]

However, the table shows that shocks to the higher-order moments of TFP alone can still explain a substantial fraction of the observed losses, especially in investment. Even without Epstein-Zin preferences and with a relatively low risk-aversion index, the model can still explain around a fourth of the observed loss in investment growth. Intuitively, in the model, when risk increases (either through the second-order or third-order moment of TFP), agents in the country shift away from domestic capital and into the internationally traded asset. This mechanism explains the drop in investment.

Role of negative skewness. One of our main findings is that negative skewness shocks play quantitatively important roles in driving business cycles. To see this, in rows labeled "second order" in Table 4, we show the fractions of observed losses explained under each calibration, but using an approximation of the solution of the model only to the second order, and thus effectively shutting down the endogenous response to the shock to the skewness of TFP. As the baseline column shows, *the reaction to skewness is substantial*: the fractions of average losses explained in the third-order rows are roughly doubling those explained in the second-order rows. Differences of comparable magnitudes are also found in the two other calibration columns.

Why does skewness matter? Intuitively, agents in our model dislike negative skewness. To see this, let $\tilde{C}_t = C_t - \theta\omega^{-1}g^{t-1}H_t^{\omega}$, the aggregate of utility from consumption and labor to the household. By the definition of household preferences, $V_t^{1-\varsigma} = (1-\beta)\tilde{C}_t^{1-\varsigma} + \beta E_t\left[V_{t+1}^{1-\gamma}\right]^{\frac{1-\varsigma}{1-\gamma}}$. Let $v_t \equiv V_t^{1-\varsigma}$ so that when $\gamma = \varsigma$ and thus Epstein-Zin preferences reduce to expected utility preferences, v_t is the usual definition of the value function for the household: $v_t = (1-\beta)\tilde{C}_t^{1-\varsigma} + \beta E_t\left[v_{t+1}^{\frac{1-\gamma}{1-\varsigma}}\right]^{\frac{1-\varsigma}{1-\gamma}}$. The third-order Taylor approximation for v_t around

[13] For example, if sectors of the economy differ not only with respect to average productivity, but also exposure to political uncertainty under unrest, we might see a reallocation of capital to relatively inefficient sectors, driving up the share of output growth loss explained. Recent work by Acemoglu et al. (2014) provides evidence that, during the Egyptian experience of the Arab Spring, firms that had closer ties to the threatened regime suffered greater losses on the Egyptian stock market than firms that did not. Exploring the macroeconomic significance of this and other *micro* risks associated with political unrest would be complementary to our analysis and is outside of the scope of this paper.

$v_{t+1} = \mu \equiv E_t[v_{t+1}]$ is:

$$\begin{aligned} v_t = & \quad (1-\beta)\tilde{C}_t^{1-\varsigma} \\ & + \beta E_t\left[v_{t+1}\right] \\ & - \beta \frac{\gamma-\varsigma}{2\mu(1-\varsigma)} Var_t\left[v_{t+1}\right] \\ & + \beta \frac{(\gamma-\varsigma)(\gamma+1-2\varsigma)}{6(\mu(1-\varsigma))^2} Skew_t\left[v_{t+1}\right] Var_t\left[v_{t+1}\right]^{3/2}. \end{aligned} \quad (2)$$

The first three terms of the continuation payoff are well-known in the literature on Epstein-Zin preferences (e.g., Colacito et al. 2013). The first term is current utility. The second is the same discounted continuation payoff that appears in non-Epstein-Zin expected utility preferences. The third term is a "correction" to expected utility that penalizes future variance of the value function as long as $\gamma > \varsigma$.[14] The fourth term is novel to a third-order approximation. Under the same assumption that $\gamma > \varsigma$ and $\varsigma < 1$, this term rewards positive skewness of the future value function and penalizes negative skewness. As γ increases, the penalties for both volatility and negative skewness increase.

The term $Skew_t\left[v_{t+1}\right] Var_t\left[v_{t+1}\right]^{3/2}$ is equal to $E_t[(v_{t+1}-\mu)^3]$, the third central moment of the value function. It shows that, for a given amount of skewness, the size of the third central moment increases in the variance. This is why skewness and variance are complementary in giving rise to precautionary motives in equilibrium.

Expression (2) is another way to see how these higher-order moments relate to a disaster risk. A disaster is an outcome on the far left tail. If variance increases, extreme events on both tails become more likely. If in addition skewness becomes more negative, the events far out on the lower tail specifically become more likely. Though we do not calculate a fourth-order approximation to this model, one can easily show that the next term in the above expansion would penalize the fourth central moment of the value function. An increase in the fourth-order moment, like an increase in negative skewness for a given second-order moment, also makes outcomes on the tails more likely. Therefore, by taking a higher-order approximation to the value function and by considering shock distributions with fat and skewed tails, we can recover some of the effects of what has been explored in the rare disaster literature.

Comparison with other studies. How do the results in Table 4 compare with other studies in the literature on the macroeconomic effects

[14] Remember, we are using a calibration where $\varsigma < 1$, so $\mu > 0$ and $\mu(1-\varsigma) > 0$.

of risk? It is well-known that increases in second-order moments lead to economic slowdowns, though the range of models in the literature is wide and none are exactly comparable with the model in this paper in terms of modeling assumptions or forcing processes. For example, while using a very different model (a closed economy with heterogeneous firms, subject to a transitory shock to the second-order moment of a composite of technology and demand, on the monthly frequency), Bloom (2009) obtains effects of risk that are of the same order of magnitude as here, i.e., doubling the standard deviation of the forcing process leads to a decline in the level of output by 2 percentage points within the first six months. For the canonical small open economy model considered here, Fernández-Villaverde et al. (2011) find that a transitory one-standard-deviation shock to second-order moment of innovations to the global interest rate (the interest rate that households in the small open economy pay on their international debt) can lead to declines in output levels in Argentina of 1.16 percentage points below steady state after sixteen quarters, or an average output growth loss of 0.29 percentage points per year, which is about 73 percent of what our baseline model predicts. Just as in Gourio's (2012) experiment with a transitory increase in the disaster probability, in our model, investment experiences the most significant decline and output contracts by a few percentage points. However, in that model, a disaster also entails some destruction of capital, so it is difficult to directly compare the two sets of numerical results.

Welfare. Finally, we evaluate the welfare loss due to the shock to the distribution of TFP. The change in the value function V_t experienced in the first period of an unrest episode corresponds to the welfare loss from facing the more negatively skewed distribution of TFP. The loss can be evaluated by considering the following counterfactual scenario: suppose that household consumption is dictated by a social planner who ensures that households enjoy labor-adjusted consumption $\bar{\tilde{C}}$ (the steady-state level of labor-adjusted consumption in the model) during each period the economy is not in unrest and $\bar{\tilde{C}}\Delta_C$, where $\Delta_C < 1$, during each period the economy is in unrest. Suppose additionally that unrest follows the same stochastic switching process as in the data and the model but there are no other sources of uncertainty to the households. The value function of the household in this scenario takes on two values: $\bar{V}$ while not in unrest, and $\bar{V}\Delta_V$, where $\Delta_V < 1$, while in unrest. The value function takes the following form:

$$\bar{V}\Delta_V = \left((1-\beta)(\bar{\tilde{C}}\Delta_C)^{1-\varsigma} + \beta \left(p(\bar{V}\Delta_V)^{1-\gamma} + (1-p)\bar{V}^{1-\gamma} \right)^{\frac{1-\varsigma}{1-\gamma}} \right)^{\frac{1}{1-\varsigma}}. \tag{3}$$

The log-linearization of the above:

$$\bar{V}^{1-\varsigma}\widehat{\Delta_V} \approx (1-\beta)\bar{\tilde{C}}^{1-\varsigma}\widehat{\Delta_C} + \beta p\bar{V}^{1-\varsigma}\widehat{\Delta_V}. \tag{4}$$

For a given $\widehat{\Delta_V}$, we can calculate the change in labor-adjusted consumption $\widehat{\Delta_C}$ that would give rise to a fall of $\widehat{\Delta_V}$ in the value function below its steady-state value for each period spent in unrest. We take $\widehat{\Delta_V}$ as calculated from our GIRF.

Our estimates imply a $\widehat{\Delta_C}$ equal to -6.1 percent. In other words, the welfare loss due to increased volatility and skewness during unrest is equal to the welfare loss if consumption were 6.1 percent lower than its steady-state value in each period of unrest. How does this number compare with those in other studies? Lucas (1987) shows that eliminating all business cycle fluctuations for a representative agent with expected-utility preferences corresponds to 0.1 percent to 0.5 percent of steady-state consumption. Dolmas (1998) finds that the same exercise under Epstein-Zin preferences yields 2 percent to 20 percent of steady-state consumption, depending on the degree of risk aversion.

3. CONCLUSION

We estimate shocks to the volatility and skewness of business cycles by exploiting the uncertainty associated with episodes of political unrest. A small open economy real business cycle model calibrated to the estimated moments from data shows that higher-order moment shocks, especially increased negative skewness, play important roles in explaining the observed average decline in economic activities. In short, the paper demonstrates the quantitative importance of time-varying skewness of shocks in the context of a small open economy real business cycle model. Our paper makes several contributions to different threads of the macroeconomic literature. In the context of real business cycle and DSGE models, the mapping from the higher-order moments of exogenous processes to moments of endogenous variables, such as the mapping studied in this paper, is relatively underexplored. While the literature has deployed a number of mechanisms (e.g., adjustment costs on investment, debt-elastic interest rates, habit in consumption, and interest rate smoothing; see Smets and Wouters 2007) to help log-linearized models better replicate the first-order and second-order moments of observed time series, it is less clear how these mechanisms affect the model's ability to match third-order moments as well. Our paper suggests it may be important to know more about the endogenous mechanisms that help or hinder matching higher-order moments of models, given that these moments could be important for the consequences of aggregate risk. Additionally, our method of accurately cal-

culating the GIRF of third-order moments may help future researchers analyze the dynamics of higher-order moments of macroeconomic aggregates in DSGE models while avoiding Monte Carlo error.

APPENDIX: A. ONLINE APPENDIX: DATA

A.1 Details of NAVCO Unrest Data

NAVCO provides detailed information on 250 nonviolent and violent mass political campaigns between 1945 and 2006. These campaigns constitute a "consensus population" of all known cases satisfying the following conditions. Each episode is a series of observable (i.e., tactics used are overt and documented), continuous (distinguishing from one-off events or revolts) mass tactics or events that mobilize nonstate actors in pursuit of a political objective. The NAVCO dataset also provides, among other information, the country, the main participating groups, the documented objective of the movement in each year of the campaign, the presence of violence in each year of the campaign, and the degree to which the movement was successful at achieving the documented objective. We focus on episodes whose objectives belong to one of the following categories:

(0) *Regime change* indicates a goal of "overthrowing the state or substantially altering state institutions to the point that it would cause a de facto shift in the regime's hold on power."

(1) *Significant institution reform* indicates a goal of "changing fundamental political structures to alleviate injustices or grant additional rights."

(2) *Policy change* indicates a goal of "changes in government policy that fall short of changes in the fundamental political structures, including changes in a state's foreign policy."

For a complete listing of NAVCO unrest episodes, see the Online Appendix C.

A.2 Estimates of Onset and Continuation Probabilities

We investigate how likely unrest is to start and how persistent it is once it starts. We establish that unrest is rare but persistent. These facts are important for understanding the economic consequences of higher-order shocks to business cycles.

Let a dummy variable U_{it} take the value of one during episodes of unrest and zero during years with no unrest, where i denotes a country and t denotes a year. We estimate both the probability of unrest onset

(i.e., the probability of unrest conditional on no unrest the previous year) and the probability of unrest continuation (i.e., the probability of unrest conditional on there being unrest in the previous year). To assess whether the probability of unrest is a function of other observable characteristics of a country, we estimate two probit models, one for onset and one for continuation. Each probit predicts $U_{it} = 1$ as a function of a constant and a vector Z_{it} of control variables, including lagged real GDP growth minus the country-specific average growth rate $\Delta Y_{i,t-1} - \frac{1}{T_i}\sum_t \Delta Y_{it}$, religious, ethnic, and linguistic fractionalization (all on a scale of 0 to 1), and income inequality (measured with the Gini coefficient). To control for region-specific factors that might influence the overall probability of a given country experiencing unrest, we include a term $\gamma_{Region(i)}$ as a region-fixed effect.[15] We do not include country-fixed effects because this would effectively exclude any country from our sample that has never experienced unrest. Instead, we want to include all countries in our sample to exploit not just variation within countries but between them as well. The fact that many countries never experience unrest is informative to estimating the probability of onset. The two probit regressions are:

$$\Pr(U_{it}|U_{it-1} = 0) = \Phi\left(\gamma_{Z0} Z_{it} + \gamma_0 + \gamma_{0,Region(i)}\right) \text{ (onset)} \tag{5}$$

$$\Pr(U_{it}|U_{it-1} = 1) = \Phi\left(\gamma_{Z1} Z_{it} + \gamma_1 + \gamma_{1,Region(i)}\right) \text{ (continuation)} \tag{6}$$

where Φ is the cumulative distribution function of the standard normal distribution.

Our baseline estimations, reported in Table 5, indicate that the onset of unrest is rare: the estimated onset probability is 1.4 percent per year. However, once it starts, unrest tends to last for several years: the estimated continuation probability in Table 6 is 83.3 percent per year. This continuation probability implies that the average duration of unrest episodes is 5.99 ($= \frac{1}{1-0.833}$).

In summary, we find that the onset of unrest is rare. But once started, unrest is persistent, leading to relatively lengthy episodes.

A.3 Political Risks Associated with Unrest

We document that the probability of large political changes increases significantly in each year of unrest. To the extent that any large po-

[15] The regions, as classified by the World Bank, are: East Asia and Pacific, Europe and Central Asia, Latin America and Caribbean, Middle East and North Africa, South Asia, and sub-Saharan Africa.

Table 5 Estimated Onset Probability

Onset	Baseline	(2)	(3)	(4)	(5)
$\Delta Y_{i,t-1} - \frac{1}{T_i}\Sigma_t \Delta Y_{it}$		-0.003	-0.004	-0.005	-0.001
		(0.01)	(0.01)	(0.02)	(0.01)
Ethnic Frac			0.500**		
			(0.24)		
Language Frac			-0.050		
			(0.21)		
Religion Frac			-0.227		
			(0.18)		
Gini				-0.014	
				(0.01)	
Europe, Central Asia					-0.095
					(0.15)
Latin America, Caribbean					0.021
					(0.15)
Middle East, North Africa					-0.005
					(0.18)
North America					*no obs.*
South Asia					0.432**
					(0.19)
Sub-Saharan Africa					0.125
					(0.14)
constant	-2.209***	-2.147***	-2.240***	-1.386***	-2.182***
	(0.03)	(0.04)	(0.11)	(0.46)	(0.11)
$\Pr(U_{i,t}\|U_{i,t-1}=0)$	0.014	0.016	0.013	0.083	0.01
	(0.00)	(0.00)	(0.00)	(0.07)	(0.00)
N	9272	5910	5357	599	5771

Notes: Probit coefficient estimates to predict onset of unrest, U_{it}, and derived probabilities. ΔY_{it} denotes real GDP growth ($= 100\ x\ (\ln Y_t - \ln Y_{t-1})$). Standard errors in parentheses. East Asia is the baseline region for the specification with region FE. *: p < 0:10. **: p < 0:05. ***: p < 0:01.

litical change entails at least a temporary disruption of the economy, an increase in the probability of disruptive events might help make sense of the increase in the left tail of the distributions of output and consumption growth documented in the next section. We estimate a series of probit regressions to predict a set of political disruptions: (1) coups, (2) positive changes in the Polity index, (3) negative changes in the Polity index, (4) large positive changes in the Polity index (greater than five points), and (5) large negative changes in the Polity index

Table 6 Estimated Continuation Probability

Continuation	Baseline	(2)	(3)	(4)	(5)
$\Delta Y_{i,t-1} - \frac{1}{T_i}\Sigma_t \Delta Y_{it}$		0.011*	0.009	-0.016	0.010
		(0.01)	(0.01)	(0.03)	(0.01)
Ethnic Frac'n			0.423		
			(0.34)		
Language Frac'n			0.406		
			(0.27)		
Religion Frac'n			-1.022**		
			(0.32)		
Gini				0.006	
				(0.02)	
Europe, Central Asia					-0.800***
					(0.26)
Latin America, Caribbean					0.317
					(0.22)
Middle East, North Africa					-0.044
					(0.28)
North America					*no obs.*
South Asia					-0.199
					(0.30)
Sub-Saharan Africa					0.007
					(0.20)
constant	0.967***	0.995***	0.974***	0.694	1.001***
	(0.06)	(0.06)	(0.18)	(0.75)	(0.18)
$\Pr(U_{i,t}\|U_{i,t-1}=1)$	0.833	0.840	0.835	0.756	0.842
	(0.01)	(0.02)	(0.04)	(0.24)	(0.04)
N	732	590	558	74	590

Notes: Probit coefficient estimates to predict onset of unrest, U_{it}, and derived probabilities. ΔY_{it} denotes real GDP growth ($= 100\ x\ (\ln Y_t - \ln Y_{t-1})$). Standard errors in parentheses. East Asia is the baseline region for the specification with region FE. *: p < 0:10. **: p < 0:05. ***: p < 0:01.

(greater than five points).[16] Each probit regression is specified as in equation (6), as a function of a constant, an indicator for current unrest, the difference between lagged real GDP growth and a country-specific average real GDP growth, and the interaction between current unrest and lagged real GDP growth. Let X_{it} be an indicator for one of the political disruptions. We estimate:

$$\Pr(X_{it}) = \Phi\left(\gamma_U U_{it} + \gamma_Z Z_{it} + \gamma_{ZU} Z_{it} U_{it} + \gamma_0\right) \tag{7}$$

[16] Data for coups come from Marshall and Marshall (2011) and data for Polity come from Marshall and Jaggers (2002).

Table 7 Estimated Probability of Political Events

	$Coup_{it}$	$\Delta Pol_{it} > 0$	$\Delta Pol_{it} < 0$	$\Delta Pol_{it} > 5$	$\Delta Pol_{it} < -5$
$U_{i,t}$	0.494***	0.802***	0.431***	0.849***	0.420**
	(0.07)	(0.07)	(0.09)	(0.10)	(0.13)
Lagged output growth†	-0.005	-0.017***	-0.005	-0.019**	-0.007
	(0.00)	(0.00)	(0.01)	(0.01)	(0.01)
$U_{i,t}$*Lagged output growth†	-0.002	0.004	0.002	0.003	0.003
	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)
constant	-1.585***	-1.696***	-1.937***	-2.397***	-2.437***
N	6500	6500	6500	6500	6500

Notes: Probit coefficient estimates to predict other political upheavals as functions of current unrest and derived probabilities. † relative to country-specific average output growth: $\Delta Y_{i,t-1} - \frac{1}{T_i}\Sigma_t \Delta Y_{it}$. Probabilities evaluated at lagged real output growth equal to country-specific average. Standard errors in parentheses. *: p < 0:10. **: p < 0:05. ***: p < 0:01.

There are a few differences between this specification and the specification of unrest onset and continuation in equation (6). First, we estimate one probit for each political disruption X_{it}. Second, in equation (6), we estimate the probits conditional on the presence of lagged unrest and the absence of lagged unrest separately. Here, we estimate one probit including both unrest and its interactions with the controls in one step. We do this to test hypotheses that the probability of each political disruption is significantly different in the presence and absence of unrest. Third, for simplicity, we include in the vector of controls Z_{it} just one control: the difference between lagged output growth and country-specific average output growth. We find that unrest is associated with increases in the probability of *all* kinds of political changes.

APPENDIX: B. ONLINE APPENDIX: MODEL DETAILS

B.1 Derivation of the Household Problem

First, we pose the problem in recursive form

$$V(K,D)^{1-\varsigma} = \max_{C,D',K',H} (1-\beta)\left(C-\theta\omega^{-1}ZH^{\omega}\right)^{1-\varsigma} + \beta E\left[V(K',D'^{1-\gamma}\right]^{\frac{1-\varsigma}{1-\gamma}}$$
$$+\lambda\left((R+1-\delta)K + WH + \frac{D'}{1+r} - D - C - K' - \frac{\phi}{2}\left(\frac{K'}{K}-g\right)^2 K\right).$$

The associated first-order conditions and envelope condition are:

$$\lambda = (1-\beta)(1-\varsigma)\left(C-\theta\omega^{-1}ZH^{\omega}\right)^{-\varsigma}$$
$$-\frac{\lambda}{1+r} = \beta(1-\varsigma)E\left[V(K',D'^{1-\gamma}\right]^{\frac{1-\varsigma}{1-\gamma}-1} E\left[V(K',D'^{-\gamma}V_D(K',D')\right]$$
$$(1-\varsigma)V(K,D)^{-\varsigma}V_D(K,D) = -\lambda$$
$$\lambda\left(1+\phi\left(\frac{K'}{K}-g\right)\right) = \beta(1-\varsigma)E\left[V(K',D'^{1-\gamma}\right]^{\frac{1-\varsigma}{1-\gamma}-1} E\left[V(K',D'^{-\gamma}V_K(K',D')\right]$$
$$(1-\varsigma)V(K,D)^{-\varsigma}V_K(K,D) = \lambda\left(R+1-\delta+\phi\left(\frac{K'}{K}-g\right)\frac{K'}{K} - \frac{\phi}{2}\left(\frac{K'}{K}-g\right)^2\right).$$

These lead to:

$$V^{1-\varsigma} = (1-\beta)\tilde{C}^{1-\varsigma} + \beta\tilde{V}^{1-\varsigma}$$
$$\tilde{V}^{1-\gamma} = E\left[V'^{1-\gamma}\right]$$
$$\tilde{C} = C - \theta\omega^{-1}ZH^{\omega}$$
$$1 = \beta E\left[\left(\frac{V'}{\tilde{V}}\right)^{\varsigma-\gamma}\left(\frac{\tilde{C}'}{\tilde{C}}\right)^{-\varsigma}(1+r)\right]$$
$$1 = \beta E\left[\left(\frac{V'}{\tilde{V}}\right)^{\varsigma-\gamma}\left(\frac{\tilde{C}'}{\tilde{C}}\right)^{-\varsigma}\left(\frac{R'+1-\delta+\phi\left(\frac{K''}{K'}-g'\right)\frac{K''}{K'}-\frac{\phi}{2}\left(\frac{K''}{K'}-g'\right)^2}{1+\phi\left(\frac{K'}{K}-g\right)}\right)\right].$$

B.2 Full Set of Equilibrium Conditions

The equilibrium conditions are (with additional variables introduced for convenience):

$$V_t^{1-\varsigma} = (1-\beta)\tilde{C}_t^{1-\varsigma} + \beta\tilde{V}_t^{1-\varsigma} \tag{8}$$

$$\tilde{V}_t^{1-\gamma} = E\left[V_{t+1}^{1-\gamma}\right] \tag{9}$$

$$\tilde{C}_t = C_t - \theta\omega^{-1}Z_{t-1}H_t^{\omega} \tag{10}$$

$$W_t = \theta Z_{t-1}H_t^{\omega-1} \tag{11}$$

$$1 = \beta E_t\left[\left(\frac{V_{t+1}}{\tilde{V}_t}\right)^{\varsigma-\gamma}\left(\frac{\tilde{C}_{t+1}}{\tilde{C}_t}\right)^{-\varsigma}(1+r_t)\right] \tag{12}$$

$$1 = \beta E_t\left[\begin{array}{c}\left(\frac{V_{t+1}}{\tilde{V}_t}\right)^{\varsigma-\gamma}\left(\frac{\tilde{C}_{t+1}}{\tilde{C}_t}\right)^{-\varsigma} \\ \left(\frac{R_{t+1}+1-\delta+\phi\left(\frac{K_{t+1}}{K_t}-g_{t+1}\right)\frac{K_{t+1}}{K_t}-\frac{\phi}{2}\left(\frac{K_{t+1}}{K_t}-g_{t+1}\right)^2}{1+\phi\left(\frac{K_t}{K_{t-1}}-g_t\right)}\right)\end{array}\right] \tag{13}$$

$$Y_t = A_tK_{t-1}^{\alpha}(Z_tH_t)^{1-\alpha} \tag{14}$$

$$W_t = (1-\alpha)A_tK_{t-1}^{\alpha}Z_t^{1-\alpha}H_t^{-\alpha} \tag{15}$$

$$R_t = \alpha A_tK_{t-1}^{\alpha-1}Z_t^{1-\alpha}H_t^{1-\alpha} \tag{16}$$

$$Y_t + \frac{D_t}{1+r_t} = D_{t-1} + C_t + I_t + \frac{\phi}{2}\left(\frac{K_t}{K_{t-1}} - g_t\right)^2 K_{t-1} \tag{17}$$

$$I_t = K_t - (1-\delta)K_{t-1} \tag{18}$$

$$r_t = r^* + \psi(e^{(\tilde{D}_t/Z_t-\bar{d})} - 1). \tag{19}$$

The equilibrium conditions, scaled ($c_t = C_t/Z_{t-1}$, $\tilde{c}_t = \tilde{C}_t/Z_{t-1}$, $h_t = H_t$, $w_t = W_t/Z_{t-1}$, $v_t = V_t/Z_{t-1}$, $\tilde{v}_t = \tilde{V}_t/Z_{t-1}$, $y_t = Y_t/Z_{t-1}$,

$k_t = K_t/Z_t$, $a_t = A_t$) and simplified:

$$v_t^{1-\varsigma} = (1-\beta)\tilde{c}_t^{1-\varsigma} + \beta\tilde{v}_t^{1-\varsigma} \tag{20}$$

$$\tilde{v}_t^{1-\gamma} = E\left[(g_t v_{t+1})^{1-\gamma}\right] \tag{21}$$

$$\tilde{c}_t = c_t - \theta\omega^{-1}h_t^{\omega} \tag{22}$$

$$w_t = \theta h_t^{\omega-1} \tag{23}$$

$$\kappa_t = \frac{k_t}{k_{t-1}} - 1 \tag{24}$$

$$1 = \beta g_t^{-\gamma} E_t\left[\left(\frac{v_{t+1}}{\tilde{v}_t}\right)^{\varsigma-\gamma}\left(\frac{\tilde{c}_{t+1}}{\tilde{c}_t}\right)^{-\varsigma}(1+r_t)\right] \tag{25}$$

$$1 = \beta g_t^{-\gamma} E_t\left[\left(\frac{v_{t+1}}{\tilde{v}_t}\right)^{\varsigma-\gamma}\left(\frac{\tilde{c}_{t+1}}{\tilde{c}_t}\right)^{-\varsigma}\left(\frac{R_{t+1}+1-\delta+\phi g_{t+1}^2(\kappa_{t+1}+\frac{1}{2}\kappa_{t+1}^2)}{1+\phi g_t\kappa_t}\right)\right] \tag{26}$$

$$y_t = a_t k_{t-1}^{\alpha}(g_t h_t)^{1-\alpha} \tag{27}$$

$$w_t h_t = (1-\alpha)y_t \tag{28}$$

$$R_t k_{t-1} = \alpha y_t \tag{29}$$

$$y_t = d_{t-1} - \frac{d_t g_t}{1+r_t} + c_t + i_t + \frac{\phi}{2}g_t^2\kappa_t^2 k_{t-1} \tag{30}$$

$$i_t = k_t g_t - (1-\delta)k_{t-1} \tag{31}$$

$$r_t = r^* + \psi(e^{(d_t-\bar{d})} - 1) \tag{32}$$

$$\log(a_{t+1}) = \rho\log(a_t) + \eta(u_t\sigma_u + (1-u_t)\sigma_q)\epsilon_{t+1} \tag{33}$$

$$\log(g_t) = \log(g_q) + \eta u_t \log(g_u) \tag{34}$$

$$\epsilon_{t+1} \sim \text{i.i.d.}N(0,1) \tag{35}$$

$$u_{t+1} \sim \text{Markov, 0 or 1 with constant transition matrix.} \tag{36}$$

Steady state at $\eta = 0$:

$$v = \left(\frac{1-\beta}{1-\beta g^{1-\varsigma}}\right)^{\frac{1}{1-\varsigma}} \tilde{c} \tag{37}$$

$$\tilde{v} = gv \tag{38}$$

$$\tilde{c} = c - \theta\omega^{-1}h^{\omega} \tag{39}$$

$$w = \theta h^{\omega-1} \tag{40}$$

$$1 = \beta g^{-\varsigma}(1+r) \tag{41}$$

$$r = R - \delta \tag{42}$$

$$y = ak^{\alpha}(gh)^{1-\alpha} \tag{43}$$

$$wh = (1-\alpha)y \tag{44}$$

$$Rk = \alpha y \tag{45}$$

$$y = d\frac{1+r-g}{1+r} + c + i \tag{46}$$

$$i = k(g - 1 + \delta) \tag{47}$$

$$r = r^{*} \tag{48}$$

$$a = 1 \tag{49}$$

$$\kappa = 0 \tag{50}$$

$$g = g_q. \tag{51}$$

B.3 Notes on Solution Method and GIRFs

To approximate the solution to equilibrium of our model, we use a higher-order perturbation method with the pruning algorithm of Andreasen et al. (2017). Because we calculate GIRFs for higher-order moments of endogenous variables, deriving analytic representations for the GIRFs, as Andreasen et al. (2017) do, would be extremely algebraically tedious. Instead, we rely on the computer algebra software Mathematica to compute these higher-order moments. This section describes our computational strategy.

The equilibrium conditions can be stated in the following form:

$$0 = E_t[F(\mathbf{y}_{t+1}, \mathbf{y}_t, \mathbf{x}_{t+1}, \mathbf{x}_t, u_{t+1}, u_t)]. \tag{52}$$

The vector of equations $\mathbf{F}$ includes all optimality conditions, constraints, and the law of motion for the exogenous process. The vector $\mathbf{y}_t$ is the vector of control variables: $[\log(y_t), \log(c_t), \log(i_t), \log(h_t), \log(r_{k,t}), \log(w_t), \log(r_t), \log(\tilde{c}_t), \kappa_t, \log(v_t), \log(\tilde{v}_t)]$. The perturbation parameter, η, is 1 in the model of interest but set to 0 at the point of approximation. The vector $\mathbf{x}_t$ is the vector of *continuous* states, including the

perturbation parameter[17]: $[\log(k_{t-1}), d_{t-1}, \log(a_t), \eta]$. u_t is the indicator for unrest, which can only take the values 0 and 1.

The solution to this model is a set of policy functions of the following form, where $\epsilon_{t+1} = \begin{bmatrix} \epsilon^u_{t+1} \\ \epsilon^q_{t+1} \end{bmatrix}$ and the two shocks ϵ^u_{t+1} and ϵ^q_{t+1} follow two i.i.d. Normal Inverse Gaussian processes, described in the text:

$$\mathbf{y}_t = \mathbf{g}(\mathbf{x}_t, u_t) \tag{53}$$

$$\mathbf{x}_{t+1} = \mathbf{h}(\mathbf{x}_t, u_t) + \eta \mathbf{S}(u_{t+1})\epsilon_{t+1}. \tag{54}$$

More specifically, for the state vector,

$$\begin{bmatrix} \log(k_t) \\ d_t \\ \log(a_{t+1}) \\ \eta \end{bmatrix} = \mathbf{h}\left(\begin{bmatrix} \log(k_{t-1}) \\ d_{t-1} \\ \log(a_t) \\ \eta \end{bmatrix}, u_t \right) + \eta \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ u_{t+1} & (1-u_{t+1}) \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \epsilon^u_{t+1} \\ \epsilon^q_{t+1} \end{bmatrix}. \tag{55}$$

At the point of approximation, the system is at a nonstochastic steady state in $\mathbf{x}_t$ and $\mathbf{y}_t$: $\mathbf{x}_t = \mathbf{x}_{ss} = [\log(k_{ss}), d_{ss}, 0, 0]$ and $\mathbf{y}_t = \mathbf{y}_{ss}$. Since the unrest and no-unrest states are completely symmetric at $\eta = 0$ by construction, the process u_t is irrelevant for the steady states of $\mathbf{x}_t$ and $\mathbf{y}_t$. Therefore, the following is true for all values of u_{t+1} and u_t:

$$0 = F(\mathbf{y}_{ss}, \mathbf{y}_{ss}, \mathbf{x}_{ss}, \mathbf{x}_{ss}, u_{t+1}, u_t). \tag{56}$$

We use a standard third-order perturbation method (e.g., Judd 1996) to construct Taylor series approximations to $\mathbf{h}(\cdot, 0)$, $\mathbf{h}(\cdot, 1)$, $\mathbf{g}(\cdot, 0)$, and $\mathbf{g}(\cdot, 1)$. Those Taylor series approximations yield the coefficients $\mathbf{h}_{0\mathbf{x}} = \frac{\partial \mathbf{h}(\mathbf{x},0)}{\partial \mathbf{x}}|_{\mathbf{x}=\mathbf{x}_{ss}}$, $\mathbf{H}_{0\mathbf{xx}} = \frac{\partial^2 \mathbf{h}(\mathbf{x},0)}{\partial \mathbf{x}^2}|_{\mathbf{x}=\mathbf{x}_{ss}}$, and $\mathbf{H}_{0\mathbf{xxx}} = \frac{\partial^3 \mathbf{h}(\mathbf{x},0)}{\partial \mathbf{x}^3}|_{\mathbf{x}=\mathbf{x}_{ss}}$, conformably reshaped:

$$\mathbf{h}(\mathbf{x}, 0) \approx \mathbf{h}_{0\mathbf{x}}\mathbf{x} + \frac{1}{2}\mathbf{H}_{0\mathbf{xx}}(\mathbf{x} \otimes \mathbf{x}) + \frac{1}{6}\mathbf{H}_{0\mathbf{xxx}}(\mathbf{x} \otimes \mathbf{x} \otimes \mathbf{x}). \tag{57}$$

This implies that the law of motion for $\mathbf{x}_t$ and $\mathbf{y}_t$ can be approximated to third order as:

$$\mathbf{x}_{t+1}|u_t, u_{t+1} = \mathbf{h}_{u_t\mathbf{x}}\mathbf{x}_t + \frac{1}{2}\mathbf{H}_{u_t\mathbf{xx}}(\mathbf{x}_t \otimes \mathbf{x}_t) + \frac{1}{6}\mathbf{H}_{u_t\mathbf{xxx}}(\mathbf{x} \otimes \mathbf{x} \otimes \mathbf{x}) + \eta \mathbf{S}_{u_{t+1}}\epsilon_{t+1} \tag{58}$$

$$\mathbf{y}_t|u_t, u_{t+1} = \mathbf{g}_{u_t\mathbf{x}}\mathbf{x}_t + \frac{1}{2}\mathbf{G}_{u_t\mathbf{xx}}(\mathbf{x}_t \otimes \mathbf{x}_t) + \frac{1}{6}\mathbf{G}_{u_t\mathbf{xxx}}(\mathbf{x}_t \otimes \mathbf{x}_t \otimes \mathbf{x}_t). \tag{59}$$

[17] We include the perturbation parameter in the definition of the state vector to simplify notation. Andreasen et al. (2017) include a brief discussion of this notation in the extensive appendix to their paper.

However, it is well-known (e.g., in Kim et al. 2008; and Den Haan and De Wind 2012) that third-order approximations like the above can have undesirable statistical properties, such as explosive simulated paths and spurious steady states. Andreasen et al. (2017) extend Kim et al. (2008) and use a pruning algorithm to eliminate these undesirable properties. The second-order pruning algorithm separates simulated components of $\mathbf{x}_t$ and $\mathbf{y}_t$ into first-order components $\mathbf{x}_t^f$ and $\mathbf{y}_t^f$, second-order components $\mathbf{x}_t^s$ and $\mathbf{y}_t^s$, and third-order components $\mathbf{x}_t^r$ and $\mathbf{y}_t^r$. The simulated quantities of interest are $\mathbf{x}_t^f + \mathbf{x}_t^s + \mathbf{x}_t^r$ and $\mathbf{y}_t^f + \mathbf{y}_t^s + \mathbf{y}_t^r$, and the components evolve linearly.

Let $\mathbf{C}_{t+1} = \mathbf{S}_{u_{t+1}}\epsilon_{t+1} + [0,0,0,1]'$. The constant vector $[0,0,0,1]'$ reflects the fact that the law of motion for the perturbation parameter is simply $\eta = 1$.

Following the approach in Andreasen et al. (2017), we have:

$$\mathbf{x}_{t+1}^f = \mathbf{h}_{\mathbf{x},u_t}\mathbf{x}_t^f + \mathbf{C}_{t+1} \tag{60}$$

$$\mathbf{x}_{t+1}^s = \mathbf{h}_{\mathbf{x},u_t}\mathbf{x}_t^s + \frac{1}{2}\mathbf{H}_{\mathbf{xx},u_t}(\mathbf{x}_t^f, \mathbf{x}_t^f) \tag{61}$$

$$\mathbf{x}_{t+1}^f \otimes \mathbf{x}_{t+1}^f = (\mathbf{h}_{\mathbf{x},u_t}\mathbf{x}_t^f + \mathbf{C}_{t+1}) \otimes (\mathbf{h}_{\mathbf{x},u_t}\mathbf{x}_t^f + \mathbf{C}_{t+1}) \tag{62}$$

$$\mathbf{x}_{t+1}^r = \mathbf{h}_{\mathbf{x},u_t}\mathbf{x}_t^r + \mathbf{H}_{\mathbf{xx},u_t}(\mathbf{x}_t^f \otimes \mathbf{x}_t^s) + \frac{1}{6}\mathbf{H}_{\mathbf{xxx},u_t}(\mathbf{x}_t^f \otimes \mathbf{x}_t^f \otimes \mathbf{x}_t^f) \tag{63}$$

$$\mathbf{x}_{t+1}^f \otimes \mathbf{x}_{t+1}^s = (\mathbf{h}_{\mathbf{x},u_t}\mathbf{x}_t^f + \mathbf{C}_{t+1}) \otimes (\mathbf{h}_{\mathbf{x},u_t}\mathbf{x}_t^s + \frac{1}{2}\mathbf{H}_{\mathbf{xx},u_t}(\mathbf{x}_t^f \otimes \mathbf{x}_t^f)) \tag{64}$$

$$\mathbf{x}_{t+1}^f \otimes \mathbf{x}_{t+1}^f \otimes \mathbf{x}_{t+1}^f = (\mathbf{h}_{\mathbf{x},u_t}\mathbf{x}_t^f + \mathbf{C}_{t+1}) \otimes (\mathbf{h}_{\mathbf{x},u_t}\mathbf{x}_t^f + \mathbf{C}_{t+1}) \otimes (\mathbf{h}_{\mathbf{x},u_t}\mathbf{x}_t^f + \mathbf{C}_{t+1}). \tag{65}$$

At this point, we deviate from the notation in Andreasen et al. (2017). Let

$$\mathbf{z}_t = \begin{bmatrix} \mathbf{x}_t^f \\ \mathbf{x}_t^s \\ \mathbf{x}_t^f \otimes \mathbf{x}_t^f \\ \mathbf{x}_t^r \\ \mathbf{x}_t^f \otimes \mathbf{x}_t^s \\ \mathbf{x}_t^f \otimes \mathbf{x}_t^f \otimes \mathbf{x}_t^f \end{bmatrix}.$$

Expanding the above, we find that

$$\mathbf{z}_{t+1} = \mathbf{A}_{u_t,u_{t+1}}(\epsilon_{t+1})\mathbf{z}_t + \mathbf{B}_{u_{t+1}}(\epsilon_{t+1}).$$

Remember that $\mathbf{C}_{t+1}$ is a function of ϵ_{t+1}.

$$\mathbf{A}_{u_t,u_{t+1}}(\epsilon_{t+1}) =$$

$$\begin{bmatrix} \mathbf{h}_{\mathbf{x},u_t} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{h}_{\mathbf{x},u_t} & \frac{1}{2}\mathbf{H}_{\mathbf{xx},u_t} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \tilde{\mathbf{H}}_{3,1} & \mathbf{0} & \mathbf{h}_{\mathbf{x},u_t} \otimes \mathbf{h}_{\mathbf{x},u_t} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{h}_{\mathbf{x},u_t} & \mathbf{H}_{\mathbf{xx},u_t} & \frac{1}{6}\mathbf{H}_{\mathbf{xxx},u_t} \\ \mathbf{0} & \mathbf{C}_{t+1} \otimes \mathbf{h}_{\mathbf{x},u_t} & \mathbf{C}_{t+1} \otimes \frac{1}{2}\mathbf{H}_{\mathbf{xx},u_t} & \mathbf{0} & \mathbf{h}_{\mathbf{x},u_t} \otimes \mathbf{h}_{\mathbf{x},u_t} & \mathbf{h}_{\mathbf{x},u_t} \otimes \frac{1}{2}\mathbf{H}_{\mathbf{xx},u_t} \\ \tilde{\mathbf{H}}_{6,1} & \mathbf{0} & \tilde{\mathbf{H}}_{6,3} & \mathbf{0} & \mathbf{0} & \mathbf{h}_{\mathbf{x},u_t} \otimes \mathbf{h}_{\mathbf{x},u_t} \otimes \mathbf{h}_{\mathbf{x},u_t} \end{bmatrix}$$

where

$$\tilde{\mathbf{H}}_{3,1} = \mathbf{h}_{\mathbf{x},u_t} \otimes \mathbf{C}_{t+1} + \mathbf{C}_{t+1} \otimes \mathbf{h}_{\mathbf{x},u_t} \tag{66}$$

$$\tilde{\mathbf{H}}_{6,1} = \mathbf{h}_{\mathbf{x},u_t} \otimes \mathbf{C}_{t+1} \otimes \mathbf{C}_{t+1} + \mathbf{C}_{t+1} \otimes \mathbf{h}_{\mathbf{x},u_t} \otimes \mathbf{C}_{t+1} + \mathbf{C}_{t+1} \otimes \mathbf{C}_{t+1} \otimes \mathbf{h}_{\mathbf{x},u_t} \tag{67}$$

$$\tilde{\mathbf{H}}_{6,3} = \mathbf{C}_{t+1} \otimes \mathbf{h}_{\mathbf{x},u_t} \otimes \mathbf{h}_{\mathbf{x},u_t} + \mathbf{h}_{\mathbf{x},u_t} \otimes \mathbf{C}_{t+1} \otimes \mathbf{h}_{\mathbf{x},u_t} + \mathbf{h}_{\mathbf{x},u_t} \otimes \mathbf{h}_{\mathbf{x},u_t} \otimes \mathbf{C}_{t+1}. \tag{68}$$

$$\mathbf{B}_{u_{t+1}}(\epsilon_{t+1}) = \begin{bmatrix} \mathbf{C}_{t+1} \\ \mathbf{0} \\ \mathbf{C}_{t+1} \otimes \mathbf{C}_{t+1} \\ \mathbf{0} \\ \mathbf{0} \\ \mathbf{0} \\ \mathbf{C}_{t+1} \otimes \mathbf{C}_{t+1} \otimes \mathbf{C}_{t+1} \end{bmatrix}.$$

Similarly, for controls $\mathbf{y}_t$, we have $(\mathbf{y}_t^f + \mathbf{y}_t^s + \mathbf{y}_t^r)|u_t = \mathbf{D}_{u_t}\mathbf{z}_t$, where

$$\mathbf{D}_{u_t} = \begin{bmatrix} \mathbf{g}_{u_t,\mathbf{x}} & \mathbf{g}_{u_t,\mathbf{x}} & \frac{1}{2}\mathbf{G}_{u_t,\mathbf{xx}} & \mathbf{g}_{\mathbf{x},u_t} & \mathbf{G}_{\mathbf{xx},u_t} & \frac{1}{6}\mathbf{G}_{\mathbf{xxx},u_t} \end{bmatrix}.$$

In this paper. we are interested in the growth rates of the controls.

$$\begin{aligned} (\Delta\mathbf{y}_{t+1})|u_t, u_{t+1} &= \left[(\mathbf{y}_{t+1}^f - \mathbf{y}_t^f) + (\mathbf{y}_{t+1}^s - \mathbf{y}_t^s) + (\mathbf{y}_{t+1}^r - \mathbf{y}_t^r)\right]|u_t, u_{t+1} \\ &= \mathbf{D}_{u_{t+1}}\mathbf{z}_{t+1} - \mathbf{D}_{u_t}\mathbf{z}_t \\ &= \mathbf{D}_{u_{t+1}}(\mathbf{A}_{u_t,u_{t+1}}(\epsilon_{t+1})\mathbf{z}_t + \mathbf{B}_{u_{t+1}}(\epsilon_{t+1})) - \mathbf{D}_{u_t}\mathbf{z}_t \\ &= (\mathbf{D}_{u_{t+1}}\mathbf{A}_{u_t,u_{t+1}}(\epsilon_{t+1}) - \mathbf{D}_{u_t})\mathbf{z}_t + \mathbf{D}_{u_{t+1}}\mathbf{B}_{u_{t+1}}(\epsilon_{t+1}) \\ &= \mathbf{A}_{u_t,u_{t+1}}^{\Delta\mathbf{y}}(\epsilon_{t+1})\mathbf{z}_t + \mathbf{B}_{u_{t+1}}^{\Delta\ \mathbf{y}}(\epsilon_{t+1}). \end{aligned}$$

To calculate the average change in the first three moments of output, investment, and consumption growth during unrest, we use the concept of GIRF from Andreasen et al. (2017) and Koop et al. (1996). In particular, we calculate the unconditional moments of all endogenous variables for two fixed paths for the unrest process. The first path is for a country that starts with no unrest and then enters into

unrest at $t = 1$ and stays there. That is, $u_t = 0\ \forall t \leq 0$ and $u_t = 1$ $\forall t \geq 1$. Denote this path for u_t as u. The second counterfactual path is one where the country never enters unrest: $u_t = 0\ \forall t$. Denote this path for u_t as $\tilde{u}$. Andreasen et al. (2017) condition on an initial value of the state vector $\mathbf{z}_0$. We instead focus on an unconditional expectation over the entire range of t to be able to arrive at a single path of moments for our exercise. The generalized IRF for the state variables $\mathbf{z}_t$ is the difference, at each point in time t, of the unconditional mean of $\mathbf{z}_t$ along the path u and the unconditional mean of $\mathbf{z}_t$ along the path $\tilde{u}$:

$$GIRF^u(\Delta\mathbf{y}_{t+1}) = E[\Delta\mathbf{y}_{t+1}|u] - E[\Delta\mathbf{y}_{t+1}|\tilde{u}]. \tag{69}$$

Andreasen et al. (2017) derive separate expressions for the evolution over time of the variances of controls. We take a different approach, which we find to be simpler, especially in dealing with third-order moments. We expand the set of objects we find a GIRF of from $\Delta\mathbf{y}_t$ to $Y_t = \begin{bmatrix} \Delta\mathbf{y}_t \\ (\Delta\mathbf{y}_t) \otimes (\Delta\mathbf{y}_t) \\ (\Delta\mathbf{y}_t) \otimes (\Delta\mathbf{y}_t) \otimes (\Delta\mathbf{y}_t) \end{bmatrix}$, so that we can compute one GIRF for all the moments of interest in one pass. For example, $vec(Var(\Delta\mathbf{y}_t)) = E[(\Delta\mathbf{y}_t) \otimes (\Delta\mathbf{y}_t)] - E[\Delta\mathbf{y}_t] \otimes E[\Delta\mathbf{y}_t]$, and the skewness of $\Delta\mathbf{y}_t$ is similarly a function of $E[(\Delta\mathbf{y}_t) \otimes (\Delta\mathbf{y}_t) \otimes (\Delta\mathbf{y}_t)]$. Using the expression $\Delta\mathbf{y}_t = \mathbf{A}^{\Delta\mathbf{y}}_{u_t,u_{t+1}}(\epsilon_{t+1})\mathbf{z}_t + \mathbf{B}^{\Delta\mathbf{y}}_{u_{t+1}}(\epsilon_{t+1})$ and expanding the Kronecker products in X_t, we have matrices $\tilde{\mathbf{A}}^{\Delta\mathbf{y}}_{u_t,u_{t+1}}(\epsilon_{t+1})$ and $\tilde{\mathbf{B}}^{\Delta\mathbf{y}}_{u_{t+1}}(\epsilon_{t+1})$ such that

$$\{X_{t+1} = \tilde{\mathbf{A}}^{\Delta\mathbf{y}}_{u_t,u_{t+1}}(\epsilon_{t+1})\{Z_t + \tilde{\mathbf{B}}^{\Delta\mathbf{y}}_{u_{t+1}}(\epsilon_{t+1}) \tag{70}$$

where $Z_t = \begin{bmatrix} \mathbf{z}_t \\ \mathbf{z}_t \otimes \mathbf{z}_t \\ \mathbf{z}_t \otimes \mathbf{z}_t \otimes \mathbf{z}_t \end{bmatrix}$.

To calculate the law of motion for Z_t, we expand the Kronecker products in the definition of Z_t using the law of motion $\mathbf{Z}_{t+1} = \mathbf{A}_{u_t,u_{t+1}}(\epsilon_{t+1})\mathbf{z}_t + \mathbf{B}_{u_{t+1}}(\epsilon_{t+1})$ and arrive at the law of motion $Z_{t+1} = \tilde{\mathbf{A}}_{u_t,u_{t+1}}(\epsilon_{t+1})Z_t| + \tilde{\mathbf{B}}_{u_{t+1}}(\epsilon_{t+1})$ for matrices $\tilde{\mathbf{A}}_{u_t,u_{t+1}}(\epsilon_{t+1})$ and $\tilde{\mathbf{B}}_{u_{t+1}}(\epsilon_{t+1})$.

For any Z_0, the independence of ϵ_{t+1} and Z_t implies for states and controls (noting that $E[\tilde{\mathbf{A}}_{u_t,u_{t+1}}(\epsilon_{t+1})|Z_0, u] = E[\tilde{\mathbf{A}}_{u_t,u_{t+1}}(\epsilon_{t+1})|u]$ and similarly for $E[\tilde{\mathbf{A}}^{\Delta\mathbf{y}}_{u_t,u_{t+1}}(\epsilon_{t+1})|Z_0, u]$, $E[\tilde{\mathbf{B}}_{u_{t+1}}(\epsilon_{t+1})|Z_0, u]$ and $E[\tilde{\mathbf{B}}^{\Delta\mathbf{y}}_{u_{t+1}}(\epsilon_{t+1})|Z_0, u]$):

$$E[Z_{t+1}|Z_0, u] = E[\tilde{\mathbf{A}}_{u_t,u_{t+1}}(\epsilon_{t+1})|u]E[Z_t|Z_0, u] + E[\tilde{\mathbf{B}}_{u_{t+1}}(\epsilon_{t+1})|u] \tag{71}$$

$$E[Y_{t+1}|Z_0, u] = E[\tilde{\mathbf{A}}^{\Delta\mathbf{y}}_{u_t,u_{t+1}}(\epsilon_{t+1})|u]E[Z_t|Z_0, u] + E[\tilde{\mathbf{B}}^{\Delta\mathbf{y}}_{u_{t+1}}(\epsilon_{t+1})|u]. \tag{72}$$

Let Z_0 be the fixed point of the law of motion for $E[Z_{t+1}Z_0, u]$, conditional on $t < 0$, in other words, conditional on no unrest at time t or $t+1$.

$$Z_0 = E[\tilde{\mathbf{A}}_{0,0}(\epsilon_{t+1})]Z_0 + E[\tilde{\mathbf{B}}_0(\epsilon_{t+1})]. \tag{73}$$

From the elements of Z_0, we can calculate the ergodic means, variances, covariances, skewnesses, and sundry third moments of all elements of $\mathbf{z}_t$ conditional on no unrest. This is the starting point of our GIRF. For $t = 1, 2, 3...$, use the laws of motion for X_t and Z_t to iterate forward, conditional on both the path u and the counterfactual path $\tilde{u}$, and use those paths to calculate the GIRF for X_t:

$$GIRF^u(X_t) = E[X_t|u] - E[X_t|\tilde{u}]. \tag{74}$$

This notation is very condensed. For example, $\tilde{\mathbf{B}}^{\Delta\mathbf{y}}_{u_{t+1}}(\epsilon_{t+1})$ is a very large matrix, with very large polynomials containing terms with order as high as ϵ^9_{t+1}. There are a large number of terms in every element of these matrices, even if expressed as Kronecker products; that is why we rely on the symbolic manipulation of Mathematica to expand these polynomials. Even after exploiting the very high degree of symmetry and redundant terms in Z_t, there are over 20,000 unique elements in that vector. Mathematica can handle these calculations very quickly, calculating the GIRFs for both states and controls in under two minutes.

1. C. ONLINE DATA APPENDIX: LIST OF NAVCO UNREST EPISODES

Table 8 List of Episodes of Mass Political Campaigns

	Country	Begin year	End year	Campaign	Target
1	Afghanistan	1978	1978	Afghans	Afghan government
2	Afghanistan	1992	1996	Taliban/ Anti-Government Forces	Afghan regime
3	Afghanistan	2001	2006	Taliban Resistance	Afghan government
4	Albania	1989	1991	Albania Anti-Communist	Communist regime
5	Algeria	1962	1963	Former Rebel Leaders	Ben Bella regime
6	Algeria	1992	2006	Islamic Salvation Front	Algerian government
7	Angola	1975	2002	UNITA	Angolan government
8	Argentina	1973	1977	ERP/ Monteneros	Argentina regime
9	Argentina	1977	1983	Argentina pro-democracy movement	Military junta
10	Argentina	1987	1987	Argentiana coup plot	Attempted coup
11	Bangladesh	1987	1990	Bangladesh Anti-Ershad	Military rule
12	Belarus	1988	1991	Belarus Anti-Communist	Communist regime
13	Belarus	2006	2006	Belarus Regime Opposition	Belarus government
14	Benin	1989	1990	Benin Anti-Communist	Communist regime
15	Bolivia	1952	1952	Bolivian Leftists	Military junta
16	Bolivia	1977	1982	Bolivian Anti-Junta	Military juntas
17	Brazil	1984	1985	Diretas ja	Military rule
18	Bulgaria	1989	1989	Bulgaria Anti-Communist	Communist regime
19	Burma	1988	2006	Karens	Burmese government
20	Burma	1988	1990	Burma pro-democracy movement	Military junta
21	Burundi	1972	1973	First Hutu Rebellion	Tutsi influence in government
22	Burundi	1988	1988	Second Hutu Rebellion	Tutsi influence in government
23	Burundi	1991	1992	Tutsi supremacists	Buyoya regime
24	Burundi	1993	2002	Third Hutu Rebellion	Power-sharing/Tutsi-dominated government
25	Cambodia	1970	1975	Khmer Rouge	Cambodian government
26	Cambodia	1978	1979	Anti-Khmer Rouge	Cambodian government
27	Cambodia	1989	1997	Second Khmer Rouge	Cambodian government
28	Chad	1968	1990	Frolinat	Chadian government
29	Chad	1994	1998	Chad rebels	Chadian regime
30	Chile	1973	1973	Pinochet-led rebels	Allende regime
31	Chile	1983	1989	Anti-Pinochet Movement	Augusto Pinochet
32	China	1956	1957	Hundred Flowers Movement	Communist regime
33	China	1966	1968	Cultural Revolution Red Guards	Anti-Maoists
34	China	1976	1979	Democracy Movement	Communist regime
35	China	1989	1989	Tiananmen	Communist regime
36	Colombia	1946	1953	Liberals of 1949	Conservative govt
37	Colombia	1964	2006	Revolutionary Armed Forces of Colombia and National Liberation Army	Colombia govt and US influence
38	Costa Rica	1948	1948	National Union Party	Calderon regime
39	Croatia	1999	2000	Croatian Institutional Reform	Semi-presidential system
40	Cuba	1956	1959	Cuban Revolution	Batista regime
41	Czechoslovakia	1989	1990	Velvet Revolution	Communist regime
42	Djibouti	1991	1994	Afar insurgency	Djibouti regime
43	Dominican Republic	1965	1965	Dominican leftists	Loyalist regime
44	Egypt	2000	2005	Kifaya	Mubarak regime
45	El Salvador	1977	1991	Salvadoran Civil Conflict	El Salvador government
46	Ethiopia	1981	1991	Tigrean People's Liberation Front	Ethiopian government
47	France	1960	1962	Pro-French Nationalists	French withdrawal from Algeria
48	Georgia	2003	2003	Rose Revolution	Shevardnadze regime
49	Ghana	1949	1950	Convention People's Party movement	British Rule
50	Ghana	2000	2000	Anti-Rawlings	Rawlings govt
51	Greece	1963	1963	Anti-Karamanlis	Karamanlis regime
52	Greece	1973	1974	Greece Anti-Military	Military rule
53	Guatemala	1954	1954	Conservative movement	Arbenz leftist regime
54	Guatemala	1961	1996	Marxist rebels (URNG)	government of Guatemala

Table 9 List of Episodes of Mass Political Campaigns (continued)

55	Guyana	1990	1990	Anti-Burnham / Hoyte	Burnham/Hoyte autocratic regime
56	Haiti	1985	1985	Anti-Duvalier	Jean Claude Duvalier
57	Hungary	1956	1956	Hungary Anti-Communist	Communist regime
58	Hungary	1989	1989	Hungary pro-dem movement	Communist regime
59	India	1967	1971	Naxalite rebellion	Indian regime
60	Indonesia	1949	1962	Darul Islam	Indonesian government
61	Indonesia	1956	1960	Indonesian leftists / Anti Sukarno	Sukarno regime
62	Indonesia	1997	1998	Anti-Suharto	Suharto rule
63	Iran	1977	1978	Iranian Revolution	Shah Reza Pahlavi
64	Iran	1981	1982	Iranian Mujahideen	Khomenei regime
65	Iran	1982	1983	KDPI	Iranian regime
66	Iraq	1959	1959	Shammar Tribe and pro-Western officers	Qassim regime
67	Iraq	1991	1991	Shiite rebellion	Hussein regime
68	Ivory Coast	2002	2005	PMIC	Incumbent regime
69	Kenya	1990	1991	Anti-Arap Moi	Daniel Arap Moi
70	Kyrgyzstan	1990	1991	Kyrgyzstan Democratic Movement	Communist regime
71	Kyrgyzstan	2005	2005	Tulip Revolution	Akayev regime
72	Laos	1960	1975	Pathet Lao	Laotian government
73	Lebanon	1958	1958	Anti-Shamun	Shamun regime
74	Lebanon	1975	1975	Lebanon leftists	Lebanese government
75	Lebanon	2005	2005	Cedar Revolution	Syrian forces
76	Liberia	1989	1990	Anti-Doe rebels	Doe regime
77	Liberia	1992	1995	NPFL & ULIMO	Johnson regime
78	Liberia	1996	1996	National patriotic forces	Liberian govt
79	Liberia	2003	2003	LURD	Taylor regime
80	Madagascar	1991	1993	Active Forces	Didier Radsiraka
81	Madagascar	2002	2002	Madagasar pro-democracy movement	Radsiraka regime
82	Malawi	1959	1959	Nyasaland African Congress	British rule
83	Malawi	1992	1993	Anti-Banda	Banda regime
84	Maldives	2003	2006	Anti-Gayoom	Maumoon Abudul Gayoom's regime
85	Mali	1990	1992	Mali Anti-Military	Military rule
86	Mexico	1987	2000	Anti-PRI	Corrupt govt
87	Mexico	2006	2006	Anti-Calderon	Calderon regime
88	Mongolia	1989	1990	Mongolian Anti-communist	Communist regime
89	Mozambique	1979	1992	Renamo	Mozambique government
90	Nepal	1990	1990	The Stir	Monarchy/Panchayat regime
91	Nepal	1996	2006	CPN-M/UPF	Nepalese government
92	Nicaragua	1978	1979	FSLN	Nicaraguan regime
93	Nicaragua	1980	1990	Contras	Sandinista regime
94	Niger	1991	1992	Niger Anti-Military	Military rule
95	Nigeria	1993	1998	Nigeria Anti-Military	Military rule
96	Oman	1964	1976	Popular Front for the Liberation of Oman and the Arab Gulf (PFLOAG)	Oman government
97	Pakistan	1968	1969	Anti-Khan	Khan regime
98	Pakistan	1983	1983	Pakistan pro-dem movement	Zia al-Huq
99	Pakistan	1994	1995	Mohajir	Pakistani government
100	Panama	1987	1989	Anti-Noriega	Noriega regime
101	Papua New Guinea	1988	1988	Bougainville Revolt	Papuan regime
102	Paraguay	1947	1947	Paraguay leftist rebellion	Morinigo regime
103	Peru	1980	1995	Sendero Luminoso (The Shining Path) Senderista Insurgency	Peruvian government
104	Peru	1996	1997	Tupac Amaru Revolutionary Movement (MRTA) - Senderista Insurgency	Peruvian government
105	Peru	2000	2000	Anti-Fujimori	Fujimori govt
106	Philippines	1946	1954	Hukbalahap Rebellion	Filipino government
107	Philippines	1972	2006	New People's Army	Filipino government
108	Philippines	1983	1986	People Power	Ferdinand Marcos
109	Philippines	2001	2001	Second People Power Movement	Estrada regime
110	Poland	1956	1956	Poznan Protests	Communist regime

Table 10 List of Episodes of Mass Political Campaigns (continued)

111	Poland	1968	1968	Poland Anti-Communist I	Communist regime
112	Poland	1970	1970	Poland Anti-Communist II	Communist regime
113	Poland	1976	1976	Poland Warsaw worker uprising	Communist regime
114	Poland	1980	1989	Solidarity	Communist regime
115	Portugal	1973	1974	Carnation Revolution	Military rule
116	Romania	1987	1989	Anti-Ceaucescu rebels	Ceacescu regime
117	Russia	1990	1991	Russia pro-dem movement	Anti-coup
118	Rwanda	1961	1964	Watusi	Hutu regime
119	Rwanda	1990	1994	Tutsi rebels	Hutu regime
120	Rwanda	1994	1994	Patriotic Front	Hutu regime and genocide
121	Senegal	2000	2000	Anti-Diouf	Diouf govt
122	Serbia	1996	2000	Anti-Milosevic	Milosevic regime
123	Sierra Leone	1991	1996	RUF	Republican government
124	Slovenia	1989	1990	Slovenia Anti-Communist	Communist regime
125	Somalia	1982	1991	Somalia clan factions; SNM	Siad Barre regime
127	South Africa	1952	1961	South Africa First Defiance Campaign	Apartheid
128	South Africa	1984	1994	South Africa Second Defiance Campaign	Apartheid
129	South Korea	1960	1960	South Korea Student Revolution	Rhee regime
130	South Korea	1979	1980	South Korea Anti-Junta	Military junta
131	South Korea	1987	1987	South Korea Anti-Military	Military government
132	Sri Lanka	1971	1971	JVP	Sri Lankan government
133	Sri Lanka	1972	1972	LTTE	Sri Lankan occupation
134	Sudan	1985	1985	Anti-Jaafar	Jaafar Nimiery
135	Sudan	1985	2005	SPLA-Garang faction	Sudanese government
136	Sudan	2003	2006	JEM/SLA	Janjaweed militia
137	Syria	1980	1982	Muslim Brotherhood	Syrian regime
138	Taiwan	1979	1985	Taiwan pro-democracy movement	Autocratic regime
139	Tajikistan	1992	1997	Popular Democratic Army (UTO)	Rakhmanov regime
140	Tanzania	1992	1992	Tanzania pro-democracy movement	Mwinyi regime
141	Thailand	1966	1981	Thai communist rebels	Thai government
142	Thailand	1973	1973	Thai student protests	Military dictatorship
143	Thailand	1992	1992	Thai pro-dem movement	Suchinda regime
144	Thailand	2005	2006	Anti-Thaksin	Thaksin regime
145	Uganda	1980	1986	National Resistance Army	Okello regime
146	Uganda	1986	2006	LRA	Museveni government
147	Ukraine	2001	2004	Orange Revolution	Kuchma regime
148	Uruguay	1963	1972	Tupamaros	Uruguay government
149	Uruguay	1984	1985	Uruguay Anti-Military	Military rule
150	Venezuela	1958	1958	Anti-Jimenez	Jimenez dictatorship
151	Venezuela	1958	1963	Armed Forces for National Liberation (FALN)	Betancourt regime
152	Yugoslavia	1968	1968	Yugoslavia student protests	Communist regime
153	Yugoslavia	1970	1971	Croatian nationalists	Yugoslav government
154	Zambia	1990	1991	Zambia Anti-Single Party	One-party rule
155	Zambia	2001	2001	Anti-Chiluba	Chiluba regime
156	Zimbabwe	1974	1979	Zimbabwe African People's Union	Smith/Muzorena regime
157	Zimbabwe	1982	1987	PF-ZAPU guerillas	Mugabe regime

REFERENCES

Acemoglu, Daron, Tarek A. Hassan, and Ahmed Tahoun. 2014. "The Power of the Street: Evidence from Egypt's Arab Spring." Working Paper 20665. Cambridge, Mass.: National Bureau of Economic Research. (November).

Acemoglu, Daron, and James A. Robinson. 2005. *Economic Origins of Dictatorship and Democracy.* Cambridge: Cambridge University Press.

Alesina, Alberto, Arnaud Devleeschauwer, William Easterly, Sergio Kurlat, and Romain Wacziarg. 2003. "Fractionalization." *Journal of Economic Growth* 8 (June):155–94.

Andersson, Jonas. 2001. "On the Normal Inverse Gaussian Stochastic Volatility Model." *Journal of Business & Economic Statistics* 19 (January): 44–54.

Andreasen, Martin M. 2012. "On the Effects of Rare Disasters and Uncertainty Shocks for Risk Premia in Non-Linear DSGE Models." *Review of Economic Dynamics* 15 (July): 295–316.

Andreasen, Martin M., Jesus Fernández-Villaverde, and Juan F. Rubio-Ramírez. 2017. "The Pruned State-Space System for Non-Linear DSGE Models: Theory and Empirical Applications." *Review of Economic Studies* 85 (January): 1–49.

Arellano, Cristina, Yan Bai, Patrick J. Kehoe. 2012. "Financial Frictions and Fluctuations in Volatility." Federal Reserve Bank of Minneapolis Research Department Staff Report 466 (Revised November 2018).

Barberis, Nicholas, and Ming Huang. 2008. "Stocks as Lotteries: The Implications of Probability Weighting for Security Prices." *American Economic Review* 98 (December): 2066–100.

Barndorff-Nielsen, Ole E. 1997. "Normal Inverse Gaussian Distributions and Stochastic Volatility Modelling." *Scandinavian Journal of Statistics* 24 (March): 1–13.

Barro, Robert J. 2006. "Rare Disasters and Asset Markets in the Twentieth Century." *Quarterly Journal of Economics* 121 (August) :823–66.

Bloom, Nicholas. 2009. "The Impact of Uncertainty Shocks." *Econometrica* 77 (May): 623–85.

Caldara, Dario, Jesus Fernandez-Villaverde, Juan F. Rubio-Ramirez, and Wen Yao. 2012. "Computing DSGE Models with Recursive Preferences and Stochastic Volatility." *Review of Economic Dynamics* 15 (April): 188–206.

Chari, V. V., Patrick J. Kehoe, and Ellen R. McGrattan. 2007. "Business Cycle Accounting." *Econometrica* 75 (May): 781–836.

Chenoweth, Erica, and Orion A. Lewis. 2013. "Unpacking Nonviolent Campaigns: Introducing the NAVCO 2.0 Dataset." *Journal of Peace Research* 50 (May): 415–23.

Chenoweth, Eerica, and Maria J. Stephan. 2011. *Why Civil Resistance Works: The Strategic Logic of Nonviolent Conflict*. New York: Columbia University Press.

Christiano, Lawrence J., Roberto Motto, and Massimo Rostagno. 2014. "Risk Shocks." *American Economic Review* 104 (January): 27–65.

Colacito, Riccardo, Max Croce, Steven Ho, and Philip Howard. 2013. "BKK the EZ Way: An International Production Economy with Recursive Preferences." Working Paper (April).

Colacito, Riccardo, and Mariano M. Croce. 2013. "International Asset Pricing with Recursive Preferences." *Journal of Finance* 68 (December): 2651–686.

Colacito, Riccardo, Eric Ghysels, Jinghan Meng, and Wasin Siwasarit. 2015. "Skewness in Expected Macro Fundamentals and the Predictability of Equity Returns: Evidence and Theory." Working Paper (August).

Den Haan, Wouter J., and Joris De Wind. 2012. "Nonlinear and Stable Perturbation-Based Approximations." *Journal of Economic Dynamics and Control* 36 (October): 1477–497.

Dolmas, Jim. (1998). "Risk Preferences and the Welfare Cost of Business Cycles." *Review of Economic Dynamics* 1 (July): 646–76.

Edmond, Chris. 2013. "Information Manipulation, Coordination, and Regime Change." *Review of Economic Studies* 80 (October): 1422–458.

Epstein, Larry G., and Stanley E. Zin. 1989. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica* 57 (July): 937–69.

Fernández-Villaverde, Jesus, Pablo Guerrón-Quintana, Juan F. Rubio-Ramírez, and Martín Uribe. 2011. "Risk Matters: The Real Effects of Volatility Shocks." *American Economic Review* 101 (October): 2530–561.

Feunou, Bruno, Mohammed R. Jahan-Parvar, and Cédric Okou. 2015. "Downside Variance Risk Premium." Bank of Canada Working Paper 2015-36 (October).

Gabaix, Xavier. 2012. "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance." *Quarterly Journal of Economics* 127 (May): 645–700.

García-Cicco, Javier, Roberto Pancrazi, and Martín Uribe. 2010. "Real Business Cycles in Emerging Countries?" *American Economic Review* 100 (December): 2510–531.

Gilchrist, Simon, Jae W. Sim, and Egon Zakrajšek. 2014. "Uncertainty, Financial Frictions, and Investment Dynamics." Working Paper 20038. Cambridge, Mass.: National Bureau of Economic Research. (April).

Gourio, François. 2012. "Disaster Risk and Business Cycles." *American Economic Review* 102 (October): 2734–766.

Gourio, François, Michael Siemer, and Adrien Verdelhan. 2013. "International Risk Cycles." *Journal of International Economics* 89 (March): 471–84.

Guvenen, Fatih, Serdar Ozkan, and Jae Song. 2014. "The Nature of Countercyclical Income Risk." *Journal of Political Economy* 122 (June): 621–60.

Judd, Kenneth L. 1996. "Approximation, Perturbation, and Projection Methods in Economic Analysis." In *Handbook of Computational Economics* Vol. 1, edited by H.M. Amman, D.A. Kendrick, and J. Rust. Amsterdam: North Holland, 509–85.

Justiniano, Alejandro, and Giorgio E. Primiceri. 2008. "The Time-Varying Volatility of Macroeconomic Fluctuations." *American Economic Review* 98 (June): 604–41.

Kim, Jinill, Sunghyun Kim, Ernst Schaumburg, and Christopher A. Sims. 2008. "Calculating and Using Second-Order Accurate Solutions of Discrete Time Dynamic Equilibrium Models." *Journal of Economic Dynamics and Control* 32 (November): 3397–414.

Koop, Gary, M. Hashem Pesaran, and Simon M. Potter. 1996. "Impulse Response Analysis in Nonlinear Multivariate Models." *Journal of Econometrics* 74 (September): 119–47.

Kuran, Timur. 1989. "Sparks and Prairie Fires: A Theory of Unanticipated Political Revolution." *Public Choice* 61 (April): 41–74.

Lucas, Robert E. 1987. *Models of Business Cycles*, Volume 26. Oxford: Basil Blackwell.

Marshall, Monty G., and Keith Jaggers. 2002. "Polity IV Project: Political Regime Characteristics and Transitions, 1800-2002."

Marshall, Monty G., and Donna Ramsey Marshall. 2011. "Coup D'etat Events, 1946-2012 Codebook." Center for Systemic Peace.

Mencía, Javier, and Enrique Sentana. 2012. "Distributional Tests in Multivariate Dynamic Models with Normal and Student-t Innovations." *Review of Economics and Statistics* 94 (February): 133–52.

Nakamura, Emi, Jón Steinsson, Robert Barro, and José Ursúa. 2013. "Crises and Recoveries in an Empirical Model of Consumption Disasters." *American Economic Journal: Macroeconomics* 5 (July): 35–74.

Rancière, Romain, Aaron Tornell, and Frank Westermann. 2008. "Systemic Crises and Growth." *Quarterly Journal of Economics* 123 (February): 359–406.

Richter, Alexander W., Nathaniel A. Throckmorton, and Todd B. Walker. 2014. "Accuracy, Speed and Robustness of Policy Function Iteration." *Computational Economics* 44 (December): 445–76.

Rudebusch, Glenn D., and Eric T. Swanson. 2012. "The Bond Premium in a DSGE Model with Long-Run Real and Nominal Risks." *American Economic Journal: Macroeconomics* 4 (January): 105–43.

Salgado, Sergio, Fatih Guvenen, and Nicholas Bloom. 2015. "Skewed Business Cycles." Working Paper (June).

Schmitt-Grohé, Stephanie, and Martín Uribe. 2003. "Closing Small Open Economy Models." *Journal of International Economics* 61 (October): 163–85.

Smets, Frank, and Rafael Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97 (June): 586–606.

Vissing-Jørgensen, Annette, and Orazio P. Attanasio. 2003. "Stock-Market Participation, Intertemporal Substitution, and Risk-Aversion." *American Economic Review* 93 (May): 383–91.

Economic Quarterly—Volume 105, Number 2—Second Quarter 2019—Pages 105–132

CDS Auctions: An Overview

Erica Paulos, Bruno Sultanum, and Elliot Tobin

A credit default swap (CDS) is a credit derivative that can be used as insurance against a reference entity's credit risk, where a reference entity is either a government or corporation that has issued debt. If a party owns equal amounts of bonds and CDSs for a particular reference entity, then the party is completely insured against a negative credit event. However, unlike insurance, it is possible to own more of the CDS protection than of the underlying bond. In this way, CDS contracts make it possible to trade on an entity's credit risk without having exposure to the entity's actual bonds.

Figure 1 summarizes how CDS contracts work. A CDS contract is a bilateral agreement between a protection seller and a protection buyer. The former is taking a short position in the CDS, while the latter is taking a long position. The protection seller compensates the protection buyer if there is a credit event with respect to any of the bonds issued by the contract's reference entity. Credit events include bankruptcy, failure to pay, and restructuring, among others. In exchange, the protection buyer makes periodic interest payments to the protection seller until the contract expires.

CDS auctions are the main settlement mechanism for CDS contracts. The auction provides a unique price for the defaulted bond, which directly impacts the amount that the protection seller needs to pay the protection buyer if a credit event occurs. In this way, CDS auctions have direct influence on payouts in the CDS market, a market that had approximately $10 trillion in contracts outstanding by the end of 2007.[1] Considering the size of the CDS market, understanding how CDS auctions function is extremely important for CDS users and regulators.

Any opinions expressed are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Richmond or the Federal Reserve System.

[1] See Aldasoro and Ehlers (2018).

Figure 1 A CDS Contract

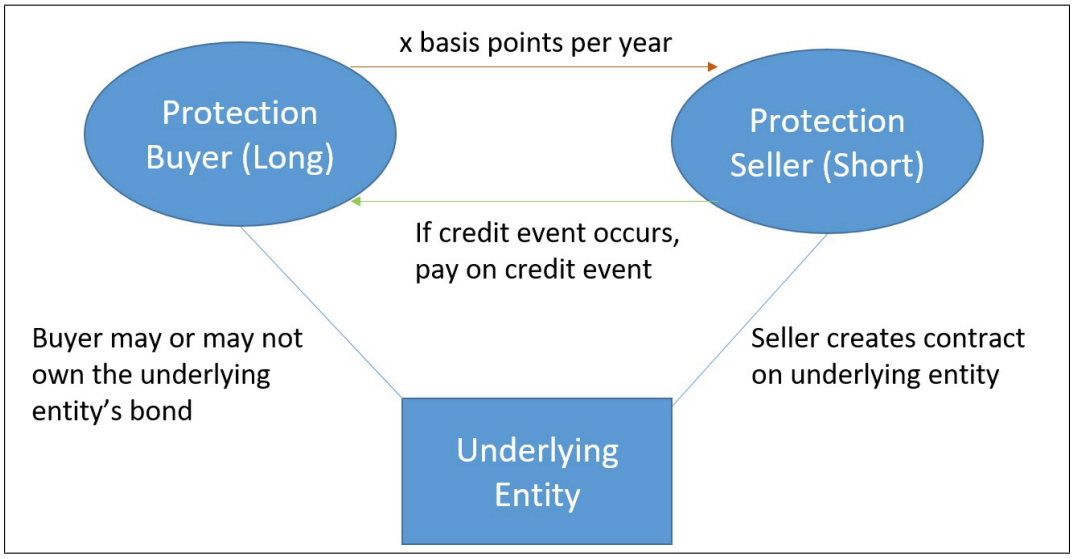

In this paper, we discuss the historical background of the CDS market, why CDS auctions were developed, and the most recent CDS literature. We then describe the auction rules and use the recent Toys R Us auction as an example. In order to illustrate frontier research on the topic, we discuss the theoretical results presented in Chernov et al. (2013) and extend their empirical findings.

Chernov et al. (2013) highlight important incentives that participants have during CDS auctions.[2] In particular, they show that dealers have an incentive to manipulate the auction price to get better terms when they settle their CDS contracts. These theoretical predictions can be empirically tested, and Chernov et al. (2013) successfully tested one of them. After extending their data to include more recent auctions, we show that their empirical results are also consistent when more recent data are included.

One difficulty that Chernov et al. (2013) face in testing some of their empirical predictions is that they do not observe dealers' CDS positions. If dealers do not actually own CDSs, they have little incentive to manipulate the auction. We further extend their work using regula-

[2] A CDS auction "participant" is anyone who wants to make a bid or offer during the auction, including dealers. "Dealers" are typically big banks that participate directly in the auction. Dealers make bids for themselves and the other auction participants.

tory data from the Depository Trust and Clearing House Corporation (DTCC) on dealers' CDS positions. Using these data, we show that some dealers have large CDS positions. This supports the theoretical findings in Chernov et al. (2013) that some dealers have an incentive to manipulate the auction price.

1. HISTORICAL BACKGROUND AND RECENT LITERATURE

JPMorgan created the first CDS contract to help manage its credit risk. After the 1989 Exxon Valdez oil spill, Exxon needed a large loan to pay for the spill's damages. Since Exxon was an important client, JPMorgan wanted to serve them but did not want the risk associated with the loan. In order to both serve Exxon and not take risk, JPMorgan made the loan to Exxon and entered into a CDS contract with the European Bank of Reconstruction and Development (EBRD). The EBRD was now responsible for covering losses resulting from Exxon defaulting on its obligation to JPMorgan. In exchange, JPMorgan agreed to pay the EBRD for the protection. In this way, CDS contracts allowed JPMorgan to transfer the loan's risk off of its books, while the EBRD was able to get exposure to the loan without having Exxon as a client.

We define a CDS's *underlying assets* as the assets (usually bonds or loans) that trigger a CDS payout if a credit event occurs with respect to them. In the example above, JPMorgan's loan to Exxon is the underlying asset.

CDS contracts allow banks to manage credit risk without trading, or even owning, the CDS's underlying bonds; however, banks are not the only possible users of CDSs. A wide variety of users, such as hedge funds, that wish to exchange credit risk trade CDSs. The initial lack of common standards for CDS contracts made it hard to trade CDSs; as a result, the market lacked liquidity. In the late 1990s, the International Swaps and Derivatives Association (ISDA) issued a set of standard credit derivatives definitions for use in connection with the ISDA Master Agreement. Combined with guidance from financial regulators, these standards helped the market grow from $632 billion in the early 2000s to $20 trillion in 2006. During the same time period, other financial and nonfinancial investors joined the CDS market, decreasing the market share of the banks. In 2000, banks accounted for 81 percent of all protection sold and 63 percent of all protection bought through CDS contracts. In 2006, banks' respective market shares fell to 59 percent and 44 percent, respectively.[3]

[3] See Mengle (2007).

In the early 1990s, most protection buyers (that is, those making a payment to transfer the credit risk) were banks carrying the underlying asset. In case of a credit event, the protection buyer could transfer the asset to the protection seller and get paid the protection amount, which was the par value of the bond. This agreement is called a *physical* settlement. Of course, if both sides agree to a payment, they could also settle the contract with a cash transfer and no asset transfer.[4] This agreement is called a *cash* settlement.

With the market growing rapidly in the late 1990s and early 2000s, issues regarding the settlement of CDS contracts emerged. By the mid-2000s, the CDS market was very different than it was in the 1990s. Investors were not necessarily carrying both the underlying asset (which at this point was mostly bonds) and the CDS. Many investors only held the CDS (*naked CDS holders*). For naked CDS holders, a physical settlement was not attractive. Protection buyers would have to buy bonds in the market in order to settle the CDS, and protection sellers would have to sell the bonds in order to cash its value. Buying and selling the bonds exposed both the long and short positions to price fluctuations. Moreover, with the volume of CDSs outstanding higher than the volume of bonds issued, the same bond had to be traded many times in the market to settle all CDS contracts. Given the over-the-counter nature of the bond market, the rush to buy the deliverable bonds artificially raised the price well above the expected recovery value. A particularly striking example of this followed the bankruptcy of Delphi Corporation in 2005. Delphi only had $2 billion in deliverable bonds for $28 billion in CDS contracts outstanding.[5]

Starting in 2005, CDS auctions were designed to solve these problems by providing a unified settlement mechanism. Investors can use CDS auctions for both physical and cash settlements. The auction identifies a price for the underlying bond, which then can be used for a cash settlement and the exchange of bonds. In a cash settlement, protection sellers pay protection buyers the par value of the underlying asset minus the auction price. Auction participants who prefer physical settlement sell their bonds in the auction and then settle in cash. The auction's cash settlement mechanism, combined with selling the bond at the auction price, replicates the payout of a physical settlement.

To see how this works, consider a participant with $100 in a particular bond and equivalent CDS protection in the same amount. Say the result of the auction is p dollars per 100 notional. If the participant submitted an order to sell his bonds in the auction, then his payout

[4] Typically, the cash transfer will be 100 minus bond price per 100 notional.

[5] See Augustin et al. (2014).

would be 100-p from the cash settlement, plus p from selling the bond, resulting in a 100-p+p=100 payout. This is the same outcome of a physical settlement, regardless of the final auction price, p.[6]

With CDS auctions determining the payouts of trillions of dollars in contracts, market participants, policymakers, and researchers began analyzing the auction. In particular, dealers and other large CDS auction participants can manipulate the price of the bond to favor their own CDS position. For example, a participant with a long (short) position in CDSs receives a higher return as the bond's auction price decreases (increases). As a result, a net buyer (seller) of protection has incentive to manipulate the auction price downward (upward). The ISDA recognized the possibility of price manipulation and designed the auction to prevent it from happening,[7] but whether the current design prevents price manipulation is a matter of theoretical and empirical research.

On the theoretical side, three recent papers have made significant progress toward understanding CDS auctions. Du and Zhu (2017) study a model of CDS auctions with restrictions to participation. They show that these restrictions bias the auction price, where bias is the difference between the asset price and its fundamental value. On the other hand, investors do not have price impact and do not bid strategically in order to manipulate the auction price because the authors consider an economy with a continuum of investors.[8] Peivandi (2017) allows for endogenous participation in the auction and solves for the optimal auction design. Peivandi shows that a CDS trader has incentives to prevent his counterparties from participating in the auction. By settling contracts in advance of the auction at better terms for his counterpart, the CDS trader can manipulate the auction price to his advantage. The better price resulting from the auction more than compensates the trader for the pre-auction settlement losses. As a result, neither full participation nor an unbiased price can be achieved. In Chernov et al. (2013), the authors consider an environment where the auction participants not only have price impact, but also have restrictions when buying/selling assets. They show that participants have

[6] Buying or selling bonds in the auction and then performing a cash settlement mimics a physical settlement regardless of the bond and CDS position of the participant.

[7] We interpret price manipulation as any participation in the auction with the intention to move the auction price of the bond away from its market price. We discuss the auction design later in the paper.

[8] Another interesting feature of Du and Zhu's (2017) model is that investors are heterogeneous in valuations, which has implications for the efficiency of the asset allocation. We find this interesting; however, we focus on the determinants of the auction price.

incentive to manipulate the auction price to profit from their existing CDS positions, resulting in a biased final price.

On the empirical side, four papers study bias in the auction price. That is, whether the auction price differs from the true fundamental price. Helwege et al. (2009) and Coudert and Gex (2010) investigate an early sample of CDS auctions, while Gupta and Sundaram (2012) and Chernov et al. (2013) investigate a more current sample. To approximate the fundamental price, these papers compare the market bond price near the time of the auction (the true price) to the final auction price. In their early sample, Helwege et al. (2009) do not find evidence of bias in the auction price. Coudert and Gex (2010), Gupta and Sundaram (2012), and Chernov et al. (2013) all find evidence of some bias in the auction price.

2. THE AUCTION

Participants in CDS auctions include nondealer participants (investors) and dealers. Since dealers are the only entities allowed to participate directly in the auction, nondealer participants must submit their requests and orders through the dealers.

CDS auctions have two stages, and each stage focuses on pricing the bonds deliverable in the auction. The auction's result is a uniform price for the auction's underlying bonds, which is the bond price used to cash settle all CDS contracts. That is, all CDS protection holders are paid 100 minus auction price per 100 notional by participants who are protection sellers. The final auction price is also used to settle all bids (offers) to buy (sell) the underlying bonds in the auction. This process is designed to mimic physical settlement, even though all CDS contracts are settled via cash.

We use the Toys R Us CDS auction as an example to clarify how the auction proceeds. We review the events leading up to their auction here. After Toys R Us filed for Chapter 11 bankruptcy on September 18, 2017, a set of fifteen dealers voted that a credit event did occur with respect to Toys R Us. This vote triggered CDS payouts and therefore a CDS auction. The Toys R Us CDS auction took place on October 11, 2017, twenty-three days after the company filed for bankruptcy.

The First Stage of the Auction

Participants may submit physical delivery requests and dealers must submit initial market quotations during the auction's first stage. Physical delivery requests are bids (offers) to buy (sell) the auction's underlying bonds at the auction's final price. The requests are used to

find the net open interest (NOI), which determines whether participants will be submitting bids or offers in the second stage. For the initial market quotations, all of the auction's dealers submit both bids and offers on a predetermined amount of the underlying bonds. These quotes help find the initial market midpoint (IMM), a price for the underlying bond that restricts second-stage orders. In this section, we detail how to calculate the NOI and IMM from the physical settlement requests and initial market quotations. Subsequently, we demonstrate how the price cap and adjustment amounts are calculated.

Physical Settlement Requests

Participants make physical settlement requests restricted by their CDS position. These physical settlement requests are used to calculate the NOI, which is carried over to the auction's second stage.

Submitting Physical Settlement Requests

All participants may submit a physical settlement request—an order to buy or sell the underlying bond at the price determined through the two-stage auction. Participants can only submit the quantity of bonds they are willing to buy or sell in these requests. They do not submit a price.

Dealers submit their physical settlement requests directly; meanwhile, nondealer participants may submit physical settlement requests through a dealer. Additionally, all participants are not allowed to submit requests above, or in the opposite direction of, their CDS position. That is, the participant can only offer (bid) to sell (buy) the underlying bonds using physical settlement if they are net buyers (sellers) of CDSs.[9] For example, a net buyer of $100 in protection can only offer to sell bonds via a physical settlement request, and the request cannot exceed $100. Note that a participant is not obligated to submit a physical settlement request, regardless of bond position.

Figure 2 displays the physical settlement requests submitted by the ten participating dealers in the Toys R Us auction. In total, four dealers submitted physical settlement requests. Bank of America, BNP Paribas, and Goldman Sachs submitted physical settlement offers. These offers could not exceed the amount of CDSs the dealers owned, indicating that these three dealers either had a long position in Toys R Us CDSs (they owned protection) or another participant who submitted a physical settlement offer through one of these dealers did. In particular, several news stories around that time mentioned that

[9] Recall from Section 1 that CDS protection holders must give the protection seller the underlying bond in exchange for par value during physical settlement. As a result, protection sellers bid to buy bonds and protection holders offer to sell bonds.

Figure 2 Toys R Us Auction Physical Settlement Requests by Dealer

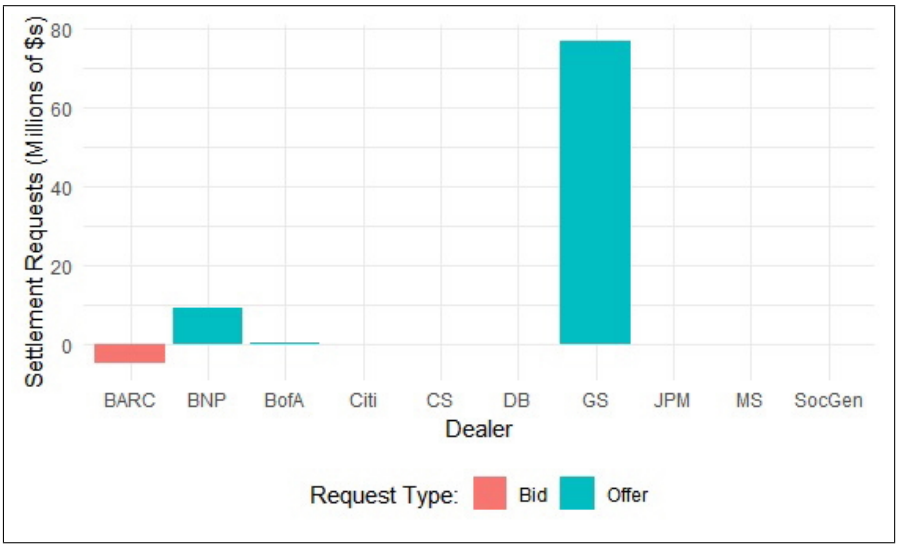

Goldman Sachs had a large long position in Toys R Us CDSs, which is consistent with their large offer to sell the underlying bond via physical delivery. Barclays was the only dealer who submitted a physical settlement bid.

Net Open Interest (NOI)

Once the physical settlement requests have been received, we calculate the NOI. Any offer during physical settlement is considered positive, while any bid is negative. We sum the physical settlement requests to find the NOI. If the sum of physical settlement offers exceeds bids (NOI is positive), then the NOI is *to sell* by the difference between the offers and bids. Likewise, the NOI is *to buy* by the difference between the bids and offers whenever physical settlement bids exceed offers (NOI is negative). If the NOI is zero, the auction ends in the first stage and the final price is set equal to the IMM. Otherwise, the NOI is taken to the second stage of the auction, where participants can bid (offer) to buy (sell) the remaining open interest depending on whether the NOI is to sell (buy).

In the Toys R Us auction, there were $5.12 million in bids and $86.292 million in offers via physical settlement requests. When offers exceed bids, the NOI is to sell and equals the amount of offers less the

amount of bids. As a result, the NOI was $81.172 million to sell,[10] so dealers submitted limit order bids in the second stage of the auction to fill the positive NOI.

Initial Market Quotations

In addition to submitting physical settlement requests during the first stage, each dealer must submit an initial market quotation, which is composed of a bid-offer pair for the underlying bonds. The first-stage quotations are used to calculate the IMM and a price cap, both of which affect the auction's second round. Additionally, either the dealer's bid or offer is carried over to the second stage, depending on the NOI. Using the Toys R Us example, we first describe the process through which dealers submit their initial market quotations; subsequently, we show how both the IMM and price cap are calculated directly from the initial quotations.

Submitting Initial Market Quotations

Dealers must submit both bids to buy and offers to sell a predetermined amount of the auction's deliverable bonds. Nondealer participants cannot submit initial quotations, even via the dealers. ISDA sets the maximum bid-ask spread prior to the auction based on asset liquidity. Any rational dealer would maximize their bid-ask spread in their quotation; therefore, each dealer's bid-ask spread should equal the maximum allowable bid-ask spread. Moreover, ISDA also sets the predetermined quotation size, which is the amount of underlying bonds for which dealers submit initial quotations.

In the Toys R Us auction, the quotation size was $2 million and the maximum bid-offer spread was 2 percent. Figure 3 displays the initial market quotations for the ten dealers who participated in the Toys R Us auction, with their bids in blue and their offers in green. As expected, each dealer's bid exceeded their offer by exactly 2 percent. Bank of America had the highest bid-offer pair (highest bid and highest offer), while Goldman Sachs and Barclays tied for the lowest (lowest bid and lowest offer).

The Initial Market Midpoint (IMM)

The dealers' initial quotations are used to calculate the IMM. Consider ordering the initial bids in descending order and the initial offers in ascending order. Then, pair the highest bid with the lowest offer, the second highest bid with the second lowest offer, and so on, until the lowest bid is paired with the highest offer. We define the bid-ask

[10] Offers are positive, and bids are negative. NOI = Offers-Bids = $86.292 million - $5.12 million = $81.172 million.

Figure 3 Toys R Us Auction Initial Quotations by Dealer

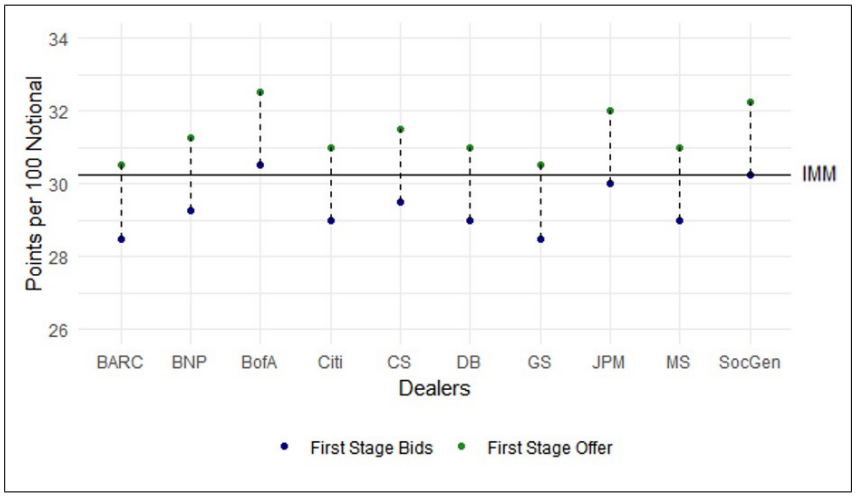

pairs created in this process as *ordered bid-offer pairs.* Any ordered bid-offer pair whose bid is greater than or equal to its offer is defined as a *crossing bid-offer pair.* In contrast, any bid-offer pair whose bid is less than its offer is a *noncrossing bid-offer pair.* The *best half* is the set of ordered bid-offer pairs whose bids are the highest among the noncrossing bid-offer pairs. In the case of an odd number of noncrossing bid-offer pairs, we round up. For instance, if there are seven noncrossing bid-offer pairs, the best half would be the four noncrossing bid-offer pairs with the highest bids (or lowest offers).

The IMM is the average of the bids and offers that make up the ordered bid-offer pairs in the best half. In other words, we first discard ordered bid-offer pairs until the ordered offer is strictly higher than the ordered bid. Then, the IMM is the average of the highest half of the remaining bids and the lowest half of the remaining offers.

Figure 4 illustrates how the IMM was calculated in the Toys R Us auction. Bids were put in descending order and offers were put in ascending order to create ordered bid-offer pairs (a bid and offer that were part of the same bid-offer pair have the same x-axis value). From the graph, a crossing bid-offer pair was identified when the bid was as high, or higher, than the offer in the ordered pair. In ordered pair 1, the highest bid (Bank of America) was equal to the lowest offer (either Barclays or Goldman Sachs) at a price of 30.5 per 100 notional. Because the bid was greater than or equal to the offer, we removed

Figure 4 Toys R Us Initial Quotation Cross to Find the IMM

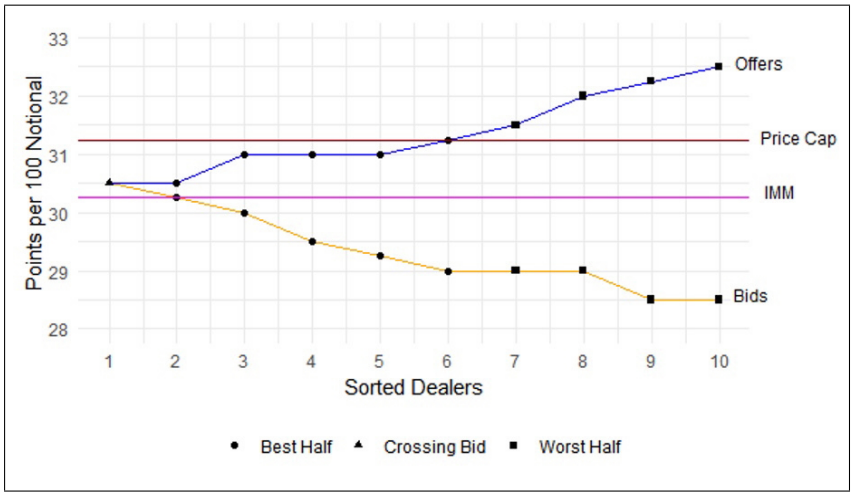

this ordered bid-offer pair (as represented by the triangle on Figure 4). There were no other crossing bid-offer pairs; therefore, only the first ordered bid-offer pair was removed. The other nine ordered bid-offer pairs made up the noncrossing pairs.

Once the crossing bids were removed, we found the best half of the remaining ordered bid-offer (noncrossing) pairs by dividing the noncrossing bid-offer pairs by two and excluding the half of pairs farthest from the cross. Since there were nine noncrossing bid-offers pairs in the Toys R Us auction, we divided by two and rounded up to get five ordered bid-offer pairs in the best half. Thus, the four ordered bid-offer pairs farthest from the cross (represented by squares in Figure 4) were also excluded from the IMM calculation. In the Toys R Us auction, there were five ordered bid-offer pairs (ten total bids and offers, represented by circles in Figure 4) comprising the best half. The IMM was the average of these ten quotations, which is 30.25 per 100 notional.

First-Stage Calculations Utilizing both the NOI and IMM

The NOI and IMM are calculated exclusively from an auction's physical settlement requests and initial quotes, respectively. In what follows, the calculation of an auction's price cap and adjustment amounts use both the NOI and IMM.

Price Cap

The price cap (floor) sets the maximum (minimum) possible auction price. When calculating a price cap, the direction of the NOI determines whether a price cap or price floor is imposed; meanwhile, the IMM determines the amount. Specifically, a price cap (floor) equal to the IMM plus (minus) half the size of the predetermined bid-offer spread is imposed if the NOI is to sell (buy). Therefore, the auction price cannot exceed (be lower than) the price cap (floor).

In the Toys R Us auction, the NOI was to sell; therefore, the auction had a price cap. The IMM was 30.25, and the bid-offer spread was 2 percent. Therefore, the second-stage price cap was 31.25 per 100 notional.[11] The brown line on Figure 4 represents the price cap in the Toys R Us auction.

Adjustment Amounts

The auction also penalizes dealers that submit initial quotes in the wrong direction of the market, via adjustment amounts. In contrast to price caps, adjustment amounts are onetime fees paid by certain dealers and therefore do not influence the auction's second stage.

The NOI determines the direction of the market, while the IMM impacts the cutoff at which dealers must pay an adjustment amount. Explicitly, if the NOI is to sell (buy), then we analyze all initial bids (offers). If a dealer's bid (offer) is higher (lower) than the IMM, then the dealer's quote is in the wrong direction of the market and they must pay the adjustment amount. It is only possible for a dealer to pay an adjustment amount if there are crossing ordered bid-offer pairs. Therefore, not every auction has a dealer that needs to pay an adjustment amount. If a dealer pays an adjustment amount, it equals the quotation amount multiplied by the amount the dealer's quote differed from the IMM.

In the Toys R Us auction, the NOI was to sell; therefore, we searched Figure 3 for any initial bids that exceeded the IMM. Bank of America's bid of 30.5 was the only initial bid (blue dots) that exceeded the IMM of 30.25 (the black horizontal line). Only Bank of America was on the wrong side of the market and paid an adjustment amount equal to the quotation amount ($2 million) multiplied by the amount their bid differed from the IMM (0.25 percent). As a result, Bank of America paid an adjustment amount of $5,000 to penalize them for being off-market.[12]

---

[11] Price Cap = IMM + (Bid-Offer Spread)/2 = 30.25 + 2/2 = 31.25.

[12] ISDA says that they round the IMM to the nearest 0.125. When calculated with greater precision, the IMM was actually 30.3125, which is equidistant to 30.25 and 30.375. We found no official rules that indicate why the IMM was rounded down instead

Figure 5 Toys R Us Initial Quotation Cross to Find the IMM

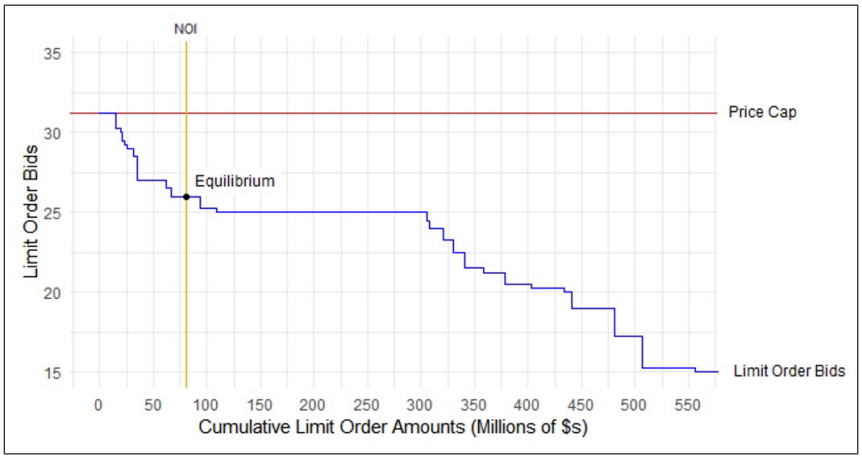

First-Stage Logistics

Participants have fifteen minutes to submit their physical settlement requests and initial quotations online via ISDA's electronic platform. Within thirty minutes of the end of this period, ISDA publishes the IMM, NOI, and any adjustment amounts on creditfixings.com. Then, participants have two to three hours to evaluate the results before the second stage begins.

The Second Stage of the Auction

In the second stage, dealers submit limit orders to fill the NOI established in the first stage to find the final auction price. This price is used to settle all outstanding CDS contracts and the auction's bond trades. In what follows, we discuss how the direction of the second stage is determined, the two ways that limit orders are submitted, and the method for determining the auction's final price. Throughout, we reference the Toys R Us auction and Figure 5.

of up. However, rounding down, as opposed to up, cost Bank of America $2,500, or 50 percent of their adjustment amount. What a rounding tragedy!

Determining the Direction of the Second Stage

The direction of the second stage depends on the NOI. If the NOI is to sell (buy), participants can only submit bids (offers) to buy (sell) the bonds. In this way, all first-round physical settlement requests that were not matched by other physical settlement requests in the first stage (NOI) are filled by limit orders in the second stage. In the Toys R Us auction, the NOI was to sell; therefore, participants submitted bids in the second stage.

Limit Orders Carried Over From First Round

The relevant side of dealers' initial quotes carry over to the second stage at the predetermined size of the initial quotation. ISDA calls these carried over quotes: *limit orders that were derived from inside markets.* If the NOI is to sell (buy), then dealers' initial bids (offers) are automatically submitted in the second stage. More specifically, as long as the relevant quote was noncrossing,[13] the relevant side of the dealer's initial quote is carried directly to the second stage. If the relevant quotation is part of a crossing bid-offer pair, then that quotation is carried forward in a different way. If this is the case and the NOI is to sell (buy), then the limit order is the minimum (maximum) of the initial bid (offer) and the IMM.

In the Toys R Us auction, dealers' initial bids for $2 million of the underlying bond were automatically carried over to the second round, since the NOI was to sell. All dealers, except Bank of America, had their initial bids carried directly to the second stage. Bank of America had the only crossing bid. For Bank of America, their mandatory second-stage bid was the minimum of the IMM (30.25) and their initial bid (30.5). As a result, Bank of America's bid derived from inside markets was 30.25 per 100 notional.

Submitting Limit Orders in the Second Stage

Any participant can submit additional limit orders by submitting a price-quantity pair in the relevant direction of the market. Unlike first-stage initial quotations, these limit orders have no predetermined size. Dealers may submit both their own limit orders and those of other participants. Each dealer can submit as many unique price-quantity pairs to buy or sell the underlying asset as they (or the participants they are

[13] Recall that a quote is noncrossing if the ordered offer is strictly greater than the ordered bid.

submitting for) want. In the Toys R Us auction, there were forty-five total limit order bids. Ten of them were derived from inside markets; meanwhile, the other thirty-five bids were submitted by participants specifically for the auction's second stage. None of the forty-five bids exceeded the price cap of 31.25 per 100 notional (brown horizontal line on Figure 5).

Some of a dealer's orders are their own, while others are those of other participants. On creditfixings.com, it is impossible to tell the orders of nondealer participants from a dealer's own orders.[14] For example, Goldman Sachs bid $81.172 million at a price of 25 in the Toys R Us auction and had smaller bids at other prices. Although it is likely that most of the smaller bids were from nondealer participants submitting bids through Goldman Sachs, it is impossible to be certain.

Finding the Auction Price

Second-stage limit orders determine the auction's bond price, which is used to cash settle the CDS contracts. The auction price is also used to settle any request to trade the underlying bond during the auction, whether through first-stage physical settlement requests or second-stage limit orders.

When the NOI is to sell (buy), we match the highest bid (lowest offer) to the amount of open interest that is equivalent to the size associated with the limit order. We continue matching the limit orders in this fashion until we match the entire NOI or run out of limit orders. In the likely case that there are sufficient limit orders to fill the NOI, the last limit order that fills the NOI is the auction price. In the unlikely scenario that there are not enough limit orders to fill the NOI, the bond price is zero (par value) when the NOI is to sell (buy). The final price is compared with the price cap. If the NOI is to sell (buy) and the final auction price is higher (lower) than the price cap (floor), then the final price is the price cap (floor)—reducing second-round price manipulation.

Dealers who submit bids (offers) that are higher (lower) than the auction price are obligated to buy (sell) the underlying bond at the final auction price along with those who submitted first-stage physical settlement requests. In this way, those who submit either a physical settlement request or a limit order in the auction exchange the relevant amount of bonds at the auction price. If there are multiple participants who made a limit order at the auction price, the amount of bonds each exchanges is proportional to the size of their limit order at the auction

[14] With the exception of the orders derived from insider markets.

price (pro rata at the margin rule). These are *partially filled orders*. All CDS contracts are then settled via cash, with the underlying bond price being the final auction price.

In the Toys R Us auction, second-stage bids were ordered from largest to smallest, creating a downward-sloping demand curve of bids (blue line in Figure 5). The NOI was a vertical supply line at $81.172 million. The point at which the downward-sloping demand curve of bids and the NOI meet is the price and quantity pair that filled the NOI (black dot in Figure 5). The equilibrium price was 26 per 100 notional, which was below the price cap of 31.25; therefore, the intersection point was the final auction price. Moreover, Barclays bid $25 million and Goldman Sachs bid $2 million at the final price of 26. When settling bond trades, Barclays and Goldman Sachs split the remaining NOI on a pro rata basis, meaning Barclays bought 25/27 of the remaining NOI from another participant who submitted a physical settlement offer in the first round.

How the Auction Prevents Manipulation

Five auction rules reduce the extent to which participants can manipulate the auction. The first three rules reduce the amount that participants can manipulate the Initial Market Midpoint (IMM) and the NOI. Rules four and five directly constrain participants' second-stage behavior.

1. Dealers' initial quotes are limited to a predetermined size and maximum spread.

2. Participants are not allowed to submit settlement requests exceeding, or in the opposite direction of, their CDS position.

3. Dealers are penalized for submitting quotes in the wrong direction of the market.

4. Based on the direction of the NOI determined in the first stage, there is a price cap (floor) on the final auction price.

5. The relevant side of all dealers' initial quotes is carried over to the second stage of the auction.

3. THE MODEL

This section and the subsequent section that together describe Chernov et al.'s (2013) model and theoretical results may be skipped. The

empirical results and main findings of this paper in Section 5 can be understood without Section 3 and 4.

In this section, we describe the basic environment (players, payoffs, etc.), the auction game, and trading frictions of CDS auctions, as discussed in Chernov et al. (2013). This section formalizes the auction rules discussed in Section 2.

Environment

Chernov et al. (2013) build their auction analysis on the work of Wilson (1979) and Back and Zender (1993). CDS auctions are two-sided auctions. There are two periods, $t = 1, 2$. The first stage of the auction happens in period $t = 1$, and the second in period $t = 2$. There is a set N of participants, from which a subset $N_d \subset N$ are dealers. Each participant $i \in N$ starts period 1 holding an endowment n_i of CDS contracts and b_i of bonds. One unit of the bond pays a value ν between zero and 100 in period $t = 1$. If ν is 100, the bond is paying the par value; if ν is zero, the bond has no residual value after default.

Holdings of CDS contracts can be zero, positive, or negative; meanwhile, bond holdings cannot be negative. Participants with positive holdings of CDS contracts are protection buyers, and participants with negative holdings are protection sellers. The net supply of CDS contracts is zero, while the supply of bonds is strictly positive. Participants have common knowledge of CDS and bond holdings.

The Auction Game

We start by describing the actions that an auction's participants can take in the first part of the auction (period $t = 1$), the second part of the auction (period $t = 2$), and the payoffs associated with such actions.

The First Stage of the Auction

In the first stage of the auction, each participant $i \in N$ submits a settlement request y_i. When y_i is positive (negative), the settlement request is an order to sell (buy) y_i units of the bond at the auction price, p^A.

The auction has restrictions on the settlement requests participants are allowed to make. A participant i with a long (short) CDS position, $n_i > 0$ ($n_i < 0$), can only submit selling (buying) orders. These orders cannot exceed their CDS position. That is, a participant with a long

position must submit a request $y_i \in [0, n_i]$, and a participant with a short position must submit a request $y_i \in [-n_i, 0]$.

In addition to settlement requests, each dealer $i \in N_d$ submits a quote π_i. Dealers must be ready to buy L units of the bond at their quoted price π_i minus a spread s or to sell L units of the bond at their quoted price π_i plus a spread s. The quotation size L and the spread s are parameters determined in advance of the auction; however, dealers do select their quote price, π_i.

After all participants submit their settlement requests and dealers submit their initial quotes, we compute the NOI and the IMM. The NOI is the sum of all settlement requests; that is, $NOI = \sum_{i \in N} y_i$. The IMM is the average of dealers' quotes, π_i, after excluding any crossing bids and offers.

The Second Stage of the Auction

There are three possibilities for the auction's second stage depending on the NOI. If the $NOI = 0$, the auction ends and the price p^A is set to the IMM. If $NOI > 0$, participants bid to buy NOI units of the bond. If $NOI < 0$, participants offer to sell $|NOI|$ units of the bond.

When $NOI > 0$, each participant $i \in N$ submits a left-continuous weakly decreasing demand schedule $x_i(p) : [0, IMM + s] \to \mathbb{R}_+$. Note that the price cap is $IMM + s$. Let $X(p) = \sum_i x_i(p)$ denote the aggregate demand for the asset at the price p. The equilibrium price is the highest price $p \in [0, IMM + s]$ such that the aggregate demand $X(p)$ matches the supply, which is the NOI. That is,

$$p^A = \max\{p | p \in [0, IMM + s] \text{ and } X(p) \geq NOI\}. \tag{1}$$

If $X(p) < NOI$ for all price $p \in [0, IMM + s]$, then p^A is set to zero. Let $q_i(p^A)$ be the asset allocation to participant $i \in N$ at the auction price p^A. The allocation is determined using a pro rata at the margin rule. Formally,

$$q_i(p^A) = x_i^+(p^A) + \frac{x_i(p^A) - x_i^+(p^A)}{X(p^A) - X^+(p^A)}(NOI - X^+(p^A)), \tag{2}$$

where $x_i^+(p^A) = \lim_{p \downarrow p^A} x_i(p)$ and $X^+(p^A) = \lim_{p \downarrow p^A} X(p)$.

When $NOI < 0$, each participant $i \in N$ submits a left-continuous weakly increasing supply function $x_i(p) : [IMM - s, 100] \to \mathbb{R}_-$; $IMM - s$ is the price floor. In this case, let $X(p) = \sum_i x_i(p)$ denote the aggregate supply for the asset at the price p. The equilibrium price is the lowest price $p \in [IMM - s, 100]$ such that the aggregate supply $X(p)$ matches the demand NOI. That is,

$$p^A = \min\{p | p \in [IMM - s, 100] \text{ and } X(p) \leq NOI\}. \tag{3}$$

If $X(p) > NOI$ for all price $p \in [IMM - s, 100]$, then p^A is set to 100. Let $q_i(p^A)$ be the bond allocation to participant $i \in N$ at the auction price p^A. As before, the allocation is determined using a pro rata at the margin rule. Formally,

$$q_i(p^A) = x_i^-(p^A) + \frac{x_i^-(p^A) - x_i(p^A)}{X^-(p^A) - X(p^A)}(NOI - X^+(p^A)), \qquad (4)$$

where $x_i^-(p^A) = \lim_{p \uparrow p^A} x_i(p)$ and $X^-(p^A) = \lim_{p \uparrow p^A} X(p)$.

Preferences

Agents are risk neutral and maximize their total payoff from the auction. The payoff of a player i given his CDS position n_i, bond holding b_i, settlement request y_i, bond allocation in the auction q_i, and final auction price p^A is

$$\Pi_i = q_i(\nu - p^A) + (n_i - y_i)(100 - p^A) + y_i(100 - \nu) + b_i\nu, \qquad (5)$$

where $\nu \in [0, 100]$ is the player's valuation for the bond. The term $q_i(\nu - p^A)$ is the player's gain from buying q_i units of bonds at the auction price p^A, $(n_i - y_i)(100 - p^A)$ is the player's gains from the cash settlement of $n_i - y_i$ units of CDS contracts. $y_i(100 - \nu)$ is the player's gains from the physical settlement of y_i units of CDS contracts, and $b_i\nu$ is the player's gains from his bond holdings.

Trading Frictions

To add realism to the environment, Chernov et al. (2013) consider two trading frictions. First, because short-selling bonds is extremely hard in practice, the authors impose the following assumption.

Assumption 1 *Each player $i \in N$ can sell at most his endowment b_i of bonds.*

Second, because some investors, such as pension funds, are not allowed to hold defaulted bonds, the authors impose the following assumption.

Assumption 2 *Only a subset of players $N_+ \subset N$, satisfying $N_+ \neq \emptyset$, can hold a positive amount of bonds after the auction.*

Solution

We analyze the model by backward induction; consequently, we start from the second stage of the auction. In the second stage, participants take all CDS positions, $\{n_i\}_i$, physical settlement requests, $\{y_i\}_i$,

and the NOI as given. Each participant i chooses his demand/supply schedule $x_i(p)$ in the auction to maximize his utility in equation (5), given the demand schedules of all players in the auction other than player i. In the first stage of the auction, participants again take all CDS positions, $\{n_i\}_i$, as given and submit settlement requests, $\{y_i\}_i$. Dealers also optimally submit their quotes $\{\pi_i\}_i$ in the first part of the auction.

All players understand that the settlement requests and dealers' quotes will determine price and quantities in the auction's second stage. This happens directly since the price in equations 1 and 3, as well as the quantities in equations 2 and 4, depend on the NOI. It also occurs indirectly since the demand/supply schedules that players submit in the auction's second stage are a function of the outcomes in the first stage. Therefore, an equilibrium is then composed of settlement requests ($\{y_i\}_i$), dealers quotes ($\{\pi_i\}_{i \in N_d}$), and demand/supply schedules ($\{x_i\}_i$) that maximize players profits given in equation (5).

4. THEORETICAL RESULTS

In this section, we describe the main theoretical results of Chernov et al. (2013). Additionally, we provide intuition for why participants have incentive to manipulate the auction price and why under/over pricing in the auction can be an equilibrium outcome. We refer the reader to Chernov et al. (2013) for the formal arguments and proofs.

Price Manipulation and the Frictionless Economy

The trading frictions implied by Assumptions 1 and 2 create limits on arbitrage. This is necessary for the result that some players manipulate the auction price to their advantage. Without market frictions, the price cannot differ from fundamentals in a meaningful way. We can conclude this result from Proposition 2 in Chernov et al. (2013), which we restate below.

Proposition 1 (2 in Chernov et al. [2013]) *Suppose there are no trading frictions, that is, Assumptions 1 and 2 are not imposed. Then, in any equilibrium, one of the following three outcomes can be realized: (1) $p \in (\nu, 100]$ and $NOI \geq 0$; (2) $p \in [0, \nu)$ and $NOI \leq 0$; and (3) $p = \nu$ and any NOI. Moreover, in all equilibria, players achieve the same expected utility as in the equilibrium with $p = \nu$.*

Proposition 1 states that, without imposing Assumptions 1 and 2, the equilibrium price for the auction, p^A, can differ from the fundamental value of ν but not in a meaningful way. In this case, every equilibrium is payoff equivalent to an equilibrium where $p^A = \nu$. Intuitively, if some equilibrium has $p^A \neq \nu$ and is not payoff equivalent to an equilibrium where $p^A = \nu$, players can gain from buying/selling bonds at the auction price. As a result, the initial allocation cannot be an equilibrium.

Price Manipulation and Trading Frictions

Once we impose trading frictions—that is, Assumptions 1 and 2—then participants can manipulate the auction price to their advantage. In this subsection, we show that manipulation is possible for the more empirically relevant case with a positive NOI ($NOI > 0$). An analogous result can be obtained when $NOI < 0$.

Proposition 2 (4 in Chernov et al. [2013]) *Suppose that there are trading frictions, that is, Assumptions 1 and 2 are imposed. Moreover, assume that*

$$\sum_{i:n_i>0} n_i \; + \sum_{i\in N_+:n_i<0} n_i \; > \; 0$$

and that for any player i who is a protection buyer, n_i satisfies

$$n_i > \frac{\sum_{j:n_j>0} n_j + \sum_{j\in N_+:n_j<0} n_j}{K+1},$$

where K is the total number of players with initial long positions ($n_i > 0$). Then, there exist a multitude of equilibria for the two-stage auction, in which $NOI > 0$ and p^A is decreasing in the NOI. In particular, there exists a subset of equilibria in which the second stage of the auction leads to a final price that is a linear function of the NOI:

$$p^A = \nu - \delta \times NOI, \tag{6}$$

where δ is defined in Chernov et al. (2013).

In general, if participants anticipate underpricing ($p^A < \nu$), protection buyers will prefer to settle in cash because they gain $100 - p^A$ from the settlement instead of $100 - \nu$, while protection sellers will prefer physical settlement because it costs them only $100 - \nu$ instead of $100 - p^A$. This would lead to a negative NOI ($NOI < 0$). However, if some of the protection sellers cannot hold bonds at the end of the auction, they will not be able to do a physical settlement. Moreover, if the protection buyers anticipate that the auction price is a negative

function of the NOI, then they will have incentive to do some physical settlement to lower the price. As a result, under some parametric restrictions, there exist equilibria where the NOI is positive and the auction price is a decreasing function of the NOI. This results in underpricing when the NOI is positive.

5. EMPIRICAL RESULTS

We replicate the test in Table 4 of Chernov et al. (2013), but we extend their data. We have their original twenty-six auctions that occurred prior to December 2011 and an additional thirteen auctions that occurred between January 2012 and December 2017.

We also look at the CDS position of dealers to verify that some dealers have long CDS positions and therefore have incentive to lower the auction price. For the dealers' CDS positions, we only have data for 2013 forward since this is when reporting CDS positions became mandatory.[15]

Data

We collect data for our regressions to replicate Chernov et al.'s (2013) results with our extended dataset. We also find dealer-level CDS positions to add to their findings.

Replication Data

To replicate Chernov et al.'s (2013) regressions, we need data for the auction's NOI, number of participants (N), aggregate CDS position ($NETCDS$), notional amount of bonds outstanding (NAB), bond price the day before the auction (p^{-1}), and the auction price (p^A).

We use creditfixings.com to find NOI, p^A, and N. They directly report each auction's NOI and p^A. Finding the number of participants is more tricky. Creditfixings.com lists all second-stage orders by dealer. Nondealers submit their bids through dealers, so we use the amount of second-stage orders and the number of dealers to estimate a lower bound for the total number of participants, following an approach proposed by Chernov et al. (2013).[16]

[15] In fact, our data have positions prior to 2013; however, since the report was not mandatory (the regulation was not in place), we are not confident the data are accurate for this period.

[16] Specifically, we get the second-stage bid/offer list, remove all orders from inside markets (dealer bids derived from the first round), and add an additional participant to the auction if a dealer submits two bids/offers at the same price (unless the dealer

The NAB represents the sum of the notional amount of bonds outstanding among each auction's deliverable obligations. Deliverable obligations are the sets of bonds that are exchanged and priced via the CDS auction at the uniform auction price, p^A. To find the NAB, we first find the Committee on Uniform Security Identification Procedures (CUSIP) number for the deliverable obligations, as listed in each auction's protocol. Then, on Fidelity's "CUSIP Look Up," we determine the initial amount of bonds offered for each deliverable obligation. Subsequently, we sum the respective initial amounts of bonds for all deliverable obligations in each auction to obtain the auction's NAB. The NAB could differ from the initial amount of bonds issued; however, the two values are unlikely to be significantly different for distressed entities. While Chernov et al. (2013) use a slightly different method for calculating the NAB, using the initial amount of bonds outstanding is a good estimate, since there is little difference between our NAB estimates and theirs for the sixteen auctions common to both samples.

We obtain an auction's $NETCDS$ from the DTCC. $NETCDS$ is at the market level. We download the weekly CDS trade data and sum the trades by week and by entity to find the $NETCDS$ position the week prior to the auction. Because Chernov et al. (2013) include auctions in which there were no $NETCDS$ data, there are fewer than thirty-seven observations of $NETCDS$ with nonmissing values.

We get the bond prices the day before the auction for the deliverable obligations (p^{-1}) using the Trade Reporting and Compliance Engine, which we access using Wharton Research Data Services. To find p^{-1}, we exclude trades below $100,000 because these trades are likely to be noninstitutional. We also remove deliverable obligations that were not the cheapest to deliver (which we define as the obligations whose prices are two standard deviations above the average price the day before the auction). Finally, p^{-1} is the average of the remaining trades.

To make sure that our data are consistent with Chernov et al. (2013), we compare the variables of interest for the twenty-six auctions that are in both datasets. Even though we do not always use the same data source as Chernov et al. (2013),[17] all the variables match well.

---

was submitting for another party, there is no reason they would submit two bids/offers at the same price). In this sense, this estimate of auction participants is a lower bound of the total possible participants.

[17] For example, Chernov et al. (2013) use Mergent data to compute the NAB and we use Fidelity's CUSIP Look Up.

Dealers' NETCDS Position

When extending Chernov et al.'s (2013) results, we find the position of individual dealers in each auction. The Dodd-Frank Wall Street Reform and Consumer Protection Act requires real-time reporting of all swap contracts to a registered swap data repository (SDR). The DTCC operates a registered SDR on CDS. The Dodd-Frank Act also requires SDRs to make all reported data available to the appropriate prudential regulators.[18] As a prudential regulator, the Federal Reserve has access to the transactions and positions involving individual parties, counterparties, or reference entities that are regulated by the Federal Reserve. Using the DTCC data, we recover the CDS position for dealers (Dealers' $NETCDS$) in fifteen auctions since 2013—a total of seventy-three observations of auction/dealer pairs or about five dealers per auction.

Regressions

In their empirical analysis, Chernov et al. (2013) investigate whether equation (6) holds in CDS auctions (Table 4). To be specific, they estimate the linear regression:

$$\frac{p^A}{p^{-1}} = \alpha + \beta \times \frac{NOI}{S} + \epsilon, \tag{7}$$

where p^A is the auction price, p^{-1} is the bond price the day before the auction, NOI is the net open interest, and S is a variable to normalize the net open interest. Since it is not clear what the normalization should be, the authors try four different specifications: no normalization ($S = 1$), number of auction participants ($S = N$), net notional amount of bond outstanding ($S = NAB$), and net CDS ($S = NETCDS$).

Table 1 depicts the empirical results in Chernov et al. (2013) and our extended dataset. The main prediction of the theory is that the β coefficient is negative. That is, p^A relative to p^{-1} is decreasing in the NOI.

In general, our results are consistent with the findings from Chernov et al. (2013). The exception is the regression in which we do not normalize the NOI. In this case, we obtain a coefficient β that is not significant, while Chernov et al. (2013) get significance at the 10 percent level. Note, however, that our p-value is 15 percent, while their

[18] See Sections 727 and 728 of the Dodd-Frank Wall Street Reform and Consumer Protection Act.

Table 1 Regression (7) for Chernov et al. (2013) and Extended Dataset

	NOI	NOI/N	NOI/NETCDS	NOI/NAB
		Chernov et al. (2013)		
α	0.90***	0.93***	0.98***	0.99***
	(20.58)	(21.59)	(17.20)	(25.29)
β	-0.07*	-3.32***	-0.41***	-0.91***
	(-1.65)	(-2.77)	(-3.07)	(-4.85)
		Using data up to December 2017		
α	0.87***	0.90***	0.92***	0.93***
	(23.89)	(26.48)	(17.99)	(23.96)
β	-0.06	-7.36***	-0.37**	-0.71***
	(-1.51)	(-3.43)	(-2.65)	(-3.26)

Notes: Sample in Chernov et al. (2013) includes twenty-six CDS auctions up to December 2011. Extended dataset includes thirty-seven CDS auctions up to December 2017. *** p<0.01, ** p<0.05, * p<0.1.

p-value is 10 percent, so our results are relatively close. Overall, we both find evidence of downward price manipulation when the NOI is positive.

Dealers' CDS Positions

The above results are consistent with the theory of price manipulation we discussed in Section 4. It also has implications for the CDS position of participants. Proposition 2 tells us that there is an equilibrium in the auction where the NOI is positive and participants with positive CDS holdings bid to lower the auction price. In our fifteen auctions, the average auction price relative the bond price the day before is 0.87, suggesting underpricing of 13 percent. Since we do observe underpricing, according to the theory, we should also have participants who are protection buyers.

Figure 6 depicts the distribution of dealers' CDS positions for the fifteen auctions for which we have dealer-level CDS positions.[19] On average, we have data for five dealers per auction; in comparison, each auction usually has ten dealers. The NOI is positive in thirteen of the fifteen auctions. After netting long and short positions for each

[19] We only have data for the auction's dealers (those directly participating in the auction). This is okay as they are likely the only ones who could have large enough positions to manipulate the auction.

Figure 6 CDS Position of Dealers in the Auction

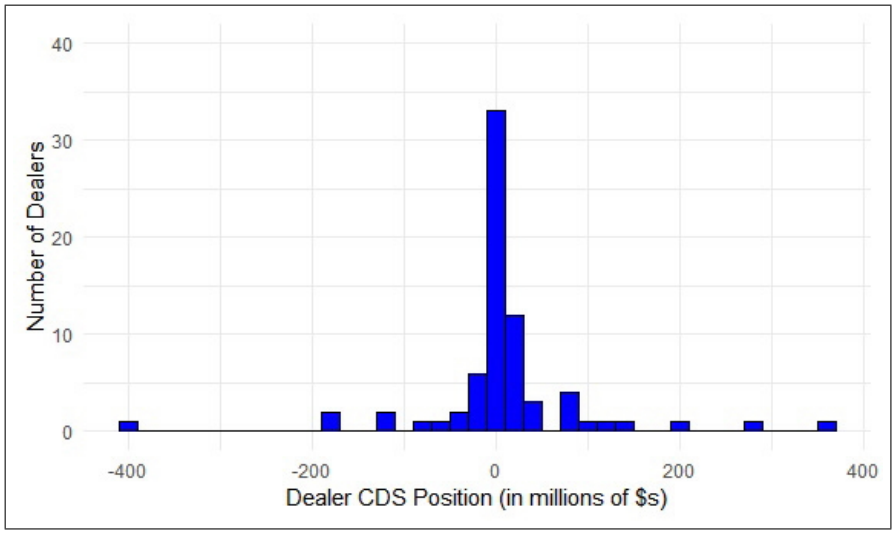

dealer, dealers own $8.23 million in protection per auction on average. Not only do dealers own protection on average, dealers also own protection in 68.5 percent of the observations. Sometimes, dealers have very large positive positions. Moreover, the CDS positions of dealers are not small in comparison to the NAB or $NETCDS$. The total notional amount of CDS holdings of dealers is 34 percent of the deliverable bonds' NAB, or 11.5 percent of the $NETCDS$—just among the dealers we observe. It seems some dealers have significant positive CDS positions and therefore have incentive to manipulate the auction price downward when the NOI is positive—supporting Chernov et al.'s (2013) empirical findings.

6. CONCLUSION

We first introduced the historical background of CDSs and CDS auctions. We then explained the auction's rules in great detail, including an example of the Toys R Us auction. These auctions are under the radar, difficult to understand, and rarely explained fully; as a result, we believe the auction details provided here will be a helpful starting point for those looking to understand CDS auctions.

After discussing three relevant CDS auction papers, we focus on Chernov et al. (2013). We provide a summary of their theoretical

model and test their empirical predictions using data through 2017. Our findings concur with theirs and indicate that CDS auction prices are being manipulated in the downward direction when the NOI is to sell.

Finally, we use regulatory data on CDS positions from the DTCC to demonstrate that dealers sometimes hold CDS positions significant enough to provide incentive to manipulate the auction. This finding provides further support for the conclusions in Chernov et al. (2013). In future work, we aim to leverage our DTCC dataset to analyze CDS auctions at the dealer holdings and bid level in a much more quantitative manner. This analysis will provide better insight into how CDS auctions are manipulated.

REFERENCES

Aldasoro, Iňaki, and Torsten Ehlers. 2018. "The Credit Default Swap Market: What a Difference a Decade Makes." *BIS Quarterly Review* (June): 1–14.

Augustin, Patrick, Marti G. Subrahmanyam, Dragon Yongjun Tang, and Sarah Qian Wang. 2014. "Credit Default Swaps: A Survey." *Foundations and Trends in Finance* 9 (December): 1–196.

Back, Kerry, and Jaime F. Zender. 2013. "Auctions of Divisible Goods: On the Rationale for the Treasury Experiment." *Review of Financial Studies* 6 (October): 733–64.

Chernov, Mikhail, Alexander S. Gorbenko, and Igor Makarov. 2013. "CDS Auctions." *Review of Financial Studies* 26 (March): 768–805.

Coudert, Virginie, and Mathieu Gex. 2010. "The Credit Default Swap Market and the Settlement of Large Defaults." *International Economics* 123 (March): 91–120.

Du, Songzi, and Haoxiang Zhu. 2017. "Are CDS Auctions Biased and Inefficient?" *Journal of Finance* 72 (December): 2589–628.

Gupta, Sudip, and Rangarajan Sundaram. 2012. "CDS Auctions and Informative Biases in CDS Recovery Rates." Working Paper (August).

Helwege, Jean, Samuel Maurer, Asani Sarkar, and Yuan Wang. 2009. "Credit Default Swap Auctions." Federal Reserve Bank of New York Staff Report 372 (May).

Mengle, David. 2007. "Credit Derivatives: An Overview." Federal Reserve Bank of Atlanta *Economic Review* 92 (December): 1–24.

Peivandi, Ahmad. 2017. "Participation and Unbiased Pricing in CDS Settlement Mechanisms." Working Paper (February).

Wilson, Robert. 1979. "Auctions of Shares." *Quarterly Journal of Economics* 93 (November): 675–89.