Science & Global Security, 23:71–100, 2015
Copyright © Taylor & Francis Group, LLC
ISSN: 0892-9882 print / 1547-7800 online
DOI: 10.1080/08929882.2015.1039431

A Bayesian Model to Assess the Size of North Korea's Uranium Enrichment Program

John E. Bistline,[1] David M. Blum,[2] Chris Rinaldi,[1] Gabriel Shields-Estrada,[3] Siegfried S. Hecker,[1,2] and M. Elisabeth Paté-Cornell[1]

[1]Department of Management Science and Engineering, Stanford University, Stanford, CA, USA
[2]Center for International Security and Cooperation, Stanford University, Stanford, CA, USA
[3]The Boston Consulting Group, Chicago, IL, USA

This article presents a model to estimate North Korea's uranium enrichment capacity and to identify probable bottlenecks for scaling up that capacity. Expert assessment is used to identify and estimate the size of key centrifuge materials and component stockpiles. Bayesian probability networks are used to characterize uncertainties in these stockpiles and a deterministic optimization model to estimate the capacity of North Korea's uranium enrichment program given the assumed components and materials constraints. A Monte Carlo simulation model is used to propagate uncertainties through the optimization model. An illustration of this approach, based on the opinions of three experts, suggests that North Korea was likely (about 80 percent chance) to have a larger uranium enrichment capacity than what was displayed to visitors to the Yongbyon nuclear complex in 2010. The three most important bottlenecks to increases in enrichment capacity are the availability of pivot bearings, maraging steel, and high-strength aluminum. The nature of the model allows it to be easily updated as new information becomes available about centrifuge materials and component stockpiles.

INTRODUCTION

In November 2010, officials from the Democratic People's Republic of Korea (North Korea) showed a team of Stanford researchers, including one of the

Received 26 May 2014; accepted 3 February 2015.
Address correspondence to John E. Bistline, Department of Management Science and Engineering, Stanford University, Huang Engineering Center, 475 Via Ortega, Stanford, CA 94305, USA. E-mail: bistline@stanford.edu
Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/gsgs.

authors, Siegfried Hecker, a modern, small uranium enrichment centrifuge facility at the Yongbyon nuclear complex. It is widely believed that the program began in the 1990s or before.[1] In the 1990s, Pakistan's A. Q. Khan sold North Korea a centrifuge starter kit with about two-dozen centrifuges and trained North Korean technicians in centrifuge operations at the Khan Research Laboratory.[2] North Korea also launched an intense, covert centrifuge materials and component procurement operation.[3] The centrifuge facility, located in a former metal fuel rod fabrication building, housed what Hecker estimated to be 2,000 centrifuges arranged in six cascades.[4] North Korea claimed at that time that the facility was operational and produced low enriched uranium (LEU) for fueling a 25–30 megawatt-electric experimental light water reactor. Pyongyang has since implied that it will reconfigure the centrifuge cascades at Yongbyon to make highly enriched uranium (HEU) instead of LEU.[5] Based on the number and type of centrifuges, the production capacity of the Yongbyon facility, if configured for weapon-grade HEU, was estimated in 2010 at roughly 40 kg/year. Overhead imagery showed that the centrifuge hall had doubled in size between 2010 and August 2013. Centrifuge facilities have a small footprint making them easy to hide, and the prospects of learning more through human intelligence in North Korea are slim. It is important to U.S. decision makers to estimate North Korea's uranium enrichment capacity and the obstacles to scaling up production.

The capacity of North Korea's uranium enrichment is measured in units of kilogram separative work units per year (kg SWU/yr.) and includes the potential to enrich uranium at current gas centrifuge facilities or in those that can be operational within the next year, including reserve centrifuges. This capacity is estimated here as a deterministic function of stockpiles of centrifuge components. This function includes an optimization of the use of the components across several types of centrifuges to maximize the enrichment capacity.

This article uses expert estimates of procurements and component stockpiles to assess indirectly the size of North Korea's uranium enrichment capacity. These estimates are from three experts, all at Stanford University at the time of the study (2014). Hecker and Englert provided an assessment of what should be considered key centrifuge components and materials, and of these, which would be the most probable bottlenecks to North Korean's centrifuge production. Hecker and Braun provided estimates of North Korean imports of centrifuge components and materials.

The next section identifies the seven key components, raw materials, and manufacturing and testing equipment for uranium enrichment centrifuges taken to be constraints on the enrichment capacity in North Korea (details are given in the Appendix). The subsequent section presents the model and how uncertainties about these stockpiles are assessed using Bayesian networks (Bayes nets) to represent the imperfect state of information about various quantities and their conditional dependencies.[6] The Bayes nets enable the derivation of probability distributions for the size of component stockpiles, which are then incorporated into a centrifuge optimization model using a Monte Carlo simulation.[7] The result, presented in the penultimate section, is a probability distribution of the enrichment capacity (in kg SWU/yr.), which is deterministically related to the uncertainties in the quantities of North Korea's stockpile of centrifuge components.

The results presented here are based on information from open sources available to the authors in 2014, and should be seen only as an illustration of this approach. To use the model for decision support in the future, the inputs will have to be updated regularly based on intelligence information and new or revised data.

Centrifuge Components

A schematic representation of a centrifuge rotor and gas centrifuge is shown in Figure 1.[8] The construction of a gas centrifuge requires several highly specific components and materials. It is assumed that all components are either indigenously produced or imported as finished products through a procurement network. In addition, many components require extensive testing before they can be used in a functioning centrifuge. Therefore, a list was developed of 30 components, raw materials, and manufacturing and testing equipment (hereafter simply "components") that are critical to the construction of a centrifuge, as shown in Table 1.[9]

In order to reduce the number of probability assessments and the computational complexity of the model, two experts evaluated this list of components and identified which ones would be the most probable bottlenecks to North Korean's centrifuge production. The question posed to the experts was whether each of the 30 components should be investigated further as a possible bottleneck or whether unlimited availability could be reasonably assumed, in which case it could be ignored. When comparing responses from the experts, a consensus emerged, for example, that the availability of raw uranium would not be a limiting factor.

Ultimately, the top seven centrifuge components were incorporated as uncertain stockpile quantities in the model (i.e., as possible bottlenecks in the North Korean centrifuge program). The corresponding random variables are thus the quantities of these components and materials available to North Korea to manufacture centrifuges. This final list included at least one component from each of the three categories of components, materials, and manufacturing/testing equipment needed in this process, namely:

- High-strength 7075 aluminum alloy
- Controller units

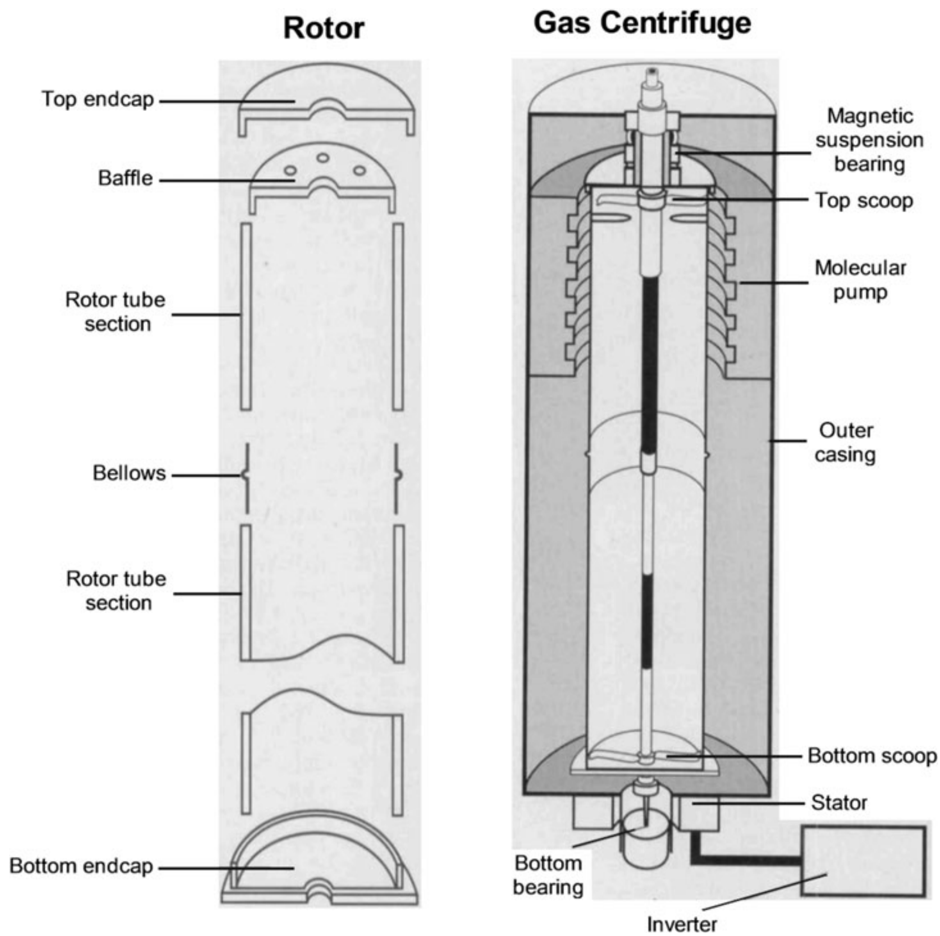

Figure 1: Schematic representations of a centrifuge rotor (left) and gas centrifuge (right). Based on Albright and Hibbs.

- Maraging steel[10] (250 and 350 grade, counted as two materials)
- Pivot bearings
- Ring magnets
- Special oils

The other 23 components shown in Table 1 are assumed to be available in such quantities that they will not be binding constraints for centrifuge production.

For special oils and ring magnets, scale independence is assumed: the North Koreans are able to obtain either an unlimited quantity of these components (by producing them indigenously or importing them) or none at all.

Table 1: Components, materials, and equipment needed to construct centrifuges

Centrifuge Components	Manufacturing/Testing	Materials
Baffles	CNC lathes	High-strength aluminum
Centrifuge housings	Controller units	Araldite potting epoxy
Coolant pipes	Flow forming machines	Maraging steel (250 and 350 grade)
End caps	Flow meters	PFPE oils
Frequency modulators	Grinders	Stainless steel
Magnetic bearings	Magnetometers	Uranium
Molecular pumps	Mass spectrometers	
Pivot bearings	Milling machines	
Bellows	Pressure transducers	
Rotor tubes	Rotor balancing equipment	
Scoops	Vacuum pumps	
Stators		

For special oils, this assumption stems from the comparatively small quantities needed to lubricate a centrifuge facility. For ring magnets, this assumption comes in part from the fact that raw samarium, which is needed to manufacture ring magnets, is widely available from China, but restrictions on manufacturing make it difficult to convince a foreign company to violate international norms.[11] Ring magnets are also similar to special oils in that their relatively small size requires fewer shipments to construct a centrifuge facility than other components, materials, and equipment on the list. In the model, private contractual arrangements are characterized with foreign companies to make ring magnets from Chinese samarium as well as acquisition of special oils from all sources as "domestic production." Within the model, "domestic production" of any material or component in Table A3 (i.e., not only scale-independent products) is treated as an all-or-nothing phenomenon. Probabilities are assigned to each of these possibilities.

Probabilistic Model

As described earlier, the model presented here determines the current probability distribution of uranium enrichment capacity in North Korea based on estimations of indigenous capacities and procurements of centrifuge-related components and materials. Estimations of the enrichment capacity are made by propagating the uncertainties about components and materials through a deterministic optimization model by Monte Carlo simulation.[12] This optimization is a maximization of the total uranium enrichment capacity (in units of kg SWU/yr.) by selecting the number of centrifuges of each design, as described below, that can be built subject to stockpile availability constraints. Figure 2 shows a diagram of the model structure. The left block presents an assessment of the probability distributions of component stockpiles. The right block pro-

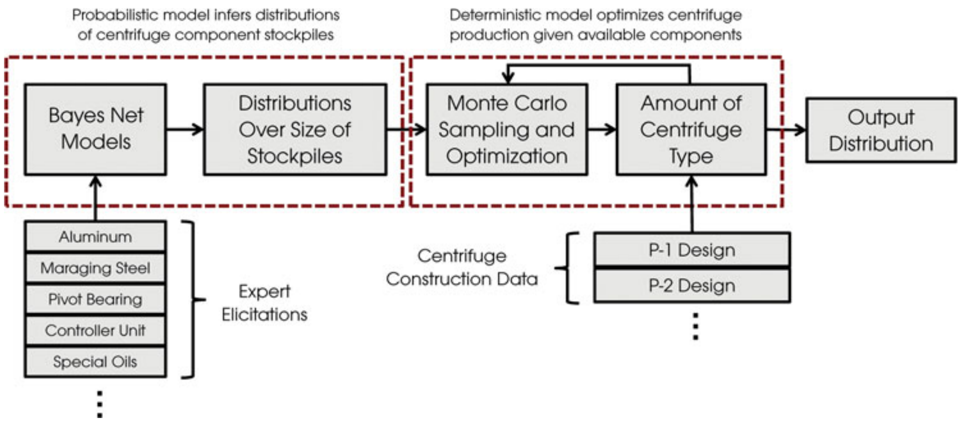

Figure 2: Diagram of model structure, from availability of materials and components to uranium enrichment capability (where P-1 and P-2 are two centrifuge types).

vides the corresponding distribution of the centrifuge capacity of North Korea at the time when the information is gathered.

For each of the uncertain stockpiles identified in the previous section, Bayes nets are created using Netica™ to determine the probability distributions of the available quantities of these components and materials. These Bayes nets provide a structured model of uncertainties (characterized by marginal and conditional distributions of random variables) and an assessment of stockpile size distributions given information from expert elicitations and open-source intelligence about known procurements.

The output distributions from these nets (available quantities) are then used as inputs to a Monte Carlo simulation model in MATLAB®. As mentioned earlier, the code samples random values from distributions of each input defining scenarios of sets of available components. In turn, the centrifuge optimization model solves for the number of P-1 and P-2 centrifuges (its decision variables) that maximize the uranium enrichment capacity, using the scenarios resulting from the Monte Carlo sampling as binding constraints. Assuming that this is actually the procedure followed by North Korea, the sampling and optimization steps are repeated until the number of runs is sufficient to construct the probability distribution of enrichment capacity. A summary of the data used in the model can be found in the Appendix.

This approach places greater emphasis on an explicit quantification of uncertainty than other assessments of North Korea's uranium enrichment program. The scenario-based approach described by Albright and Walrond evaluates production under a number of alternative assumptions about North Korea's nuclear program and its intentions.[13] Their analysis focuses on central estimates (means) of current and future stocks of weapon-grade uranium and plutonium. Although their report defines upper and lower credible bounds, their model does not attach probabilities to these various scenarios. That report also assumes the existence of a small-scale pilot plant in all but one scenario and does not account for uncertainties about the scope and total enrichment capacity at other facilities. Therefore, the resulting estimates, although useful, do not quantify the enrichment capacity based on disaggregated data about centrifuge components and materials as is done in the analysis presented here. They also do not include uncertainties about potential stockpiles.[14]

Centrifuge Optimization Model

The distributions obtained through the Bayes nets are used as inputs to a Monte Carlo simulation, as shown in Figure 3. The random values from Monte Carlo sampling are inputs to the centrifuge optimization model, which determines the optimal number of centrifuges of each type that can be constructed given the available stockpiles of materials yielded by each random sample. These model parameters include the quantities of items necessary to construct centrifuges of different types. The optimization module then determines the corresponding uranium enrichment capacity. The model considers three primary centrifuges types based on Pakistani designs from the Khan Research Laboratory: the P-1 design (with aluminum rotor), the P-2 design (with 350-grade maraging steel rotor), and a modified P-2 design (with 250-grade maraging steel rotor). This optimization framework allows identification of bottlenecks in production through binding constraints and calculates the corresponding shadow prices of these resources constraints.[15] Additional details regarding the mathematical structure of the centrifuge optimization model can be found in the Appendix.

Since the separative performances of the North Korean centrifuges (i.e., their uranium enrichment capacity) are uncertain, their values for the three centrifuge designs are treated as random variables. The model links centrifuge design parameters (e.g., rotor length) and material properties (e.g., tensile strength) to rotor speed and maximum separate work. Assuming a linear-thermal gradient profile, separative work can be calculated through the following equation:[16]

$$\delta U = \frac{V^2 Z}{33{,}000} e_E \tag{1}$$

where δU is the separative work (in kg SWU per year), V is the peripheral velocity (in meters per second), Z is the rotor length (in meters), and e_E is the dimensionless experimental efficiency.[17]

The design parameters that influence separative work for each centrifuge type (experimental efficiency, peripheral speed) are characterized as random variables with triangular distributions, which are then used as inputs in the

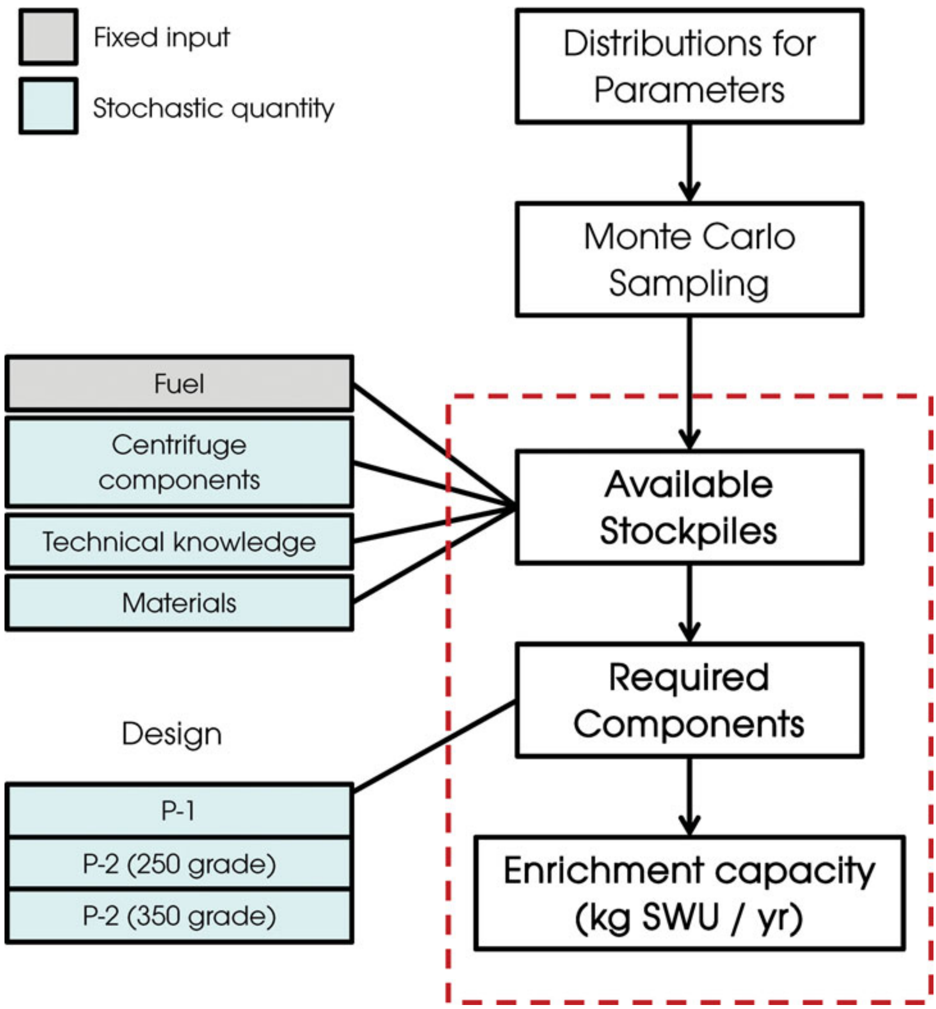

Figure 3: Schematic of the centrifuge optimization model.

Monte Carlo simulation model. The modes of these distributions come from the performance estimates provided by the North Koreans,[18] and the upper and lower limits are estimated based on the different centrifuge designs.[19] Table 2 specifies these triangular distributions.

Bayesian Network (Bayes Net) Model

To obtain the probability distributions of quantities of materials and components, a Bayes net was constructed in Netica™ to represent the dependencies and relationships among the model's random variables. As mentioned earlier, a limited number of expert opinions were used in conjunction with proba-

Table 2: Design parameters influencing separative capacity for each of three centrifuge types

Design Parameter	Centrifuge	Lower Limit	Mode	Upper Limit
Experimental efficiency	P-1	8.0%	14.5%	32.0%
	P-2 (250 grade)	20.0%	40.0%	57.0%
	P-2 (350 grade)	25.0%	60.0%	77.0%
Peripheral speed (m/s)	P-1	320	333	362
	P-2 (250 grade)	370	390	450
	P-2 (350 grade)	460	480	520

bility models to obtain the distributions required as input to the deterministic optimization model. In addition to base distributions, the Bayes nets allow the inclusion of new pieces of information in the probabilistic model by updating the distributions of the available quantities of centrifuge components.

Figure 4 represents an influence diagram that is part of the overall Bayes net. It shows that North Korea can acquire centrifuge-related materials and components through two avenues. The first is to manufacture centrifuge components in the country (indigenous capability), and the second is to acquire components and materials through a global procurement network. As an example, this influence diagram involves the (uncertain) amounts of material and components provided to North Korea by China. The model includes Europe, Russia, China, Iran, and Pakistan as potential suppliers. China is the focus for the assistance and detection sensitivities in the analysis because the probability of import from other countries is greatly reduced at this time, whereas China still presents a real potential pathway. It should be noted that this pathway is not intended to represent official Chinese government channels but the rapidly growing Chinese industrial sector.

To have the indigenous capability to produce a centrifuge component, the probabilistic model describes how North Korea must have both the *technical knowledge*[20] and *manufacturing capacity*.[21] Once North Korea has this local capability to produce a sufficient quantity of a component, that particular component will not limit its ability to construct centrifuges. If North Korea cannot manufacture a specific component, it must acquire it through its procurement network. Therefore, it must first acquire the *assistance* of another country.[22] Once it has that assistance, the model includes an assessment of the quantity of components or materials that North Korea receives from that country.

Hecker and Braun were consulted to elicit their degrees of belief about the amounts of the components and materials that North Korea has received from various sources. They provided the 10th and 90th percentiles of cumulative distribution functions (CDFs) representing these beliefs. These assessments were then fitted to a gamma distribution,[23] and that distribution was entered into the Bayes net. Similar gamma distributions were encoded for all components

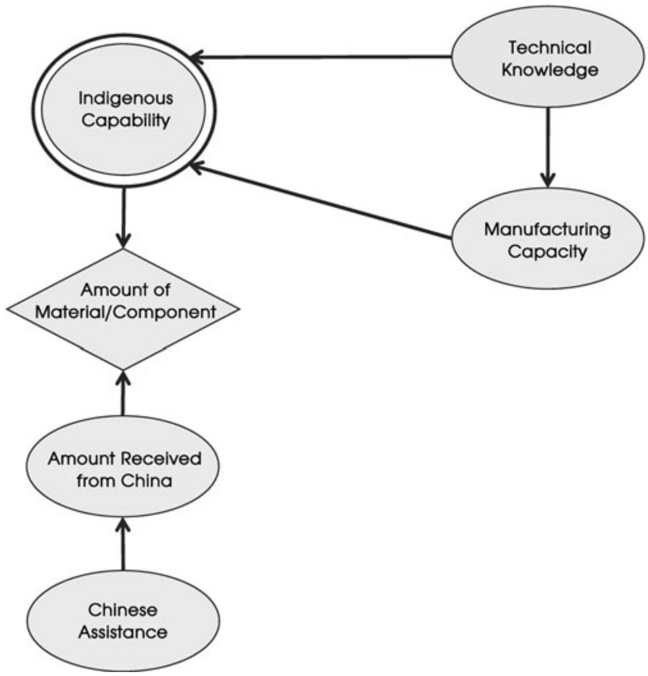

Figure 4: Influence diagram representing the (uncertain) amounts of material and components provided indigenously or through Chinese assistance.

and materials and for all possible countries of origin. Currently, only five regions are considered as potential origin (each component is assumed to have a single possible origin): Europe, Russia, China, Iran, and Pakistan. Other countries, however, could easily be added to the model, and different materials and components could have different origins.

Dependencies were included between assistance nodes for maraging steel and high-strength aluminum. The probabilistic model thus includes the fact that knowledge of one country's assistance in supplying high-strength aluminum influences the probabilities that the same country also provides assistance for the acquisition of maraging steel. Similarly, dependencies were included between technical knowledge nodes regarding maraging steel and pivot bearings, reflecting the notion that information about North Korea's technical knowledge of maraging steel production influences the probability that it has the technical knowledge of manufacturing pivot bearings. Other dependencies could be included in the probabilistic model, but computer memory requirements increase considerably as dependencies are added. Perhaps more critical, the number of dependencies increases drastically the complexity of expert opinion elicitation.

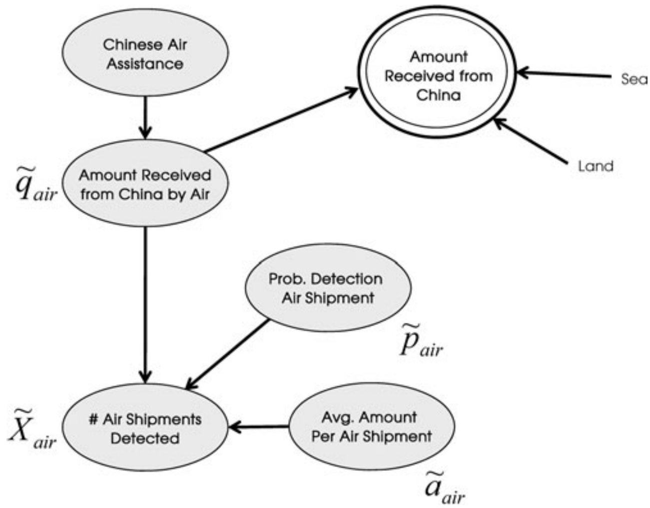

Figure 5: Detection sub-model: Bayes net representing the amount of equipment received from China.

In addition to such dependencies, a probabilistic detection sub-model was developed, also as a Netica^{TM} influence diagram. This sub-model can update distributions if, in the future, the international community detects specific instances of smuggling centrifuge-related material and components. Figure 5 shows a Bayes net used to infer the amount of a particular component imported from one country given one or more detections. The detection sub-model is discussed in the Appendix.

RESULTS

This model yields illustrative estimates of the probability distribution of North Korea's uranium enrichment capacity. It should be noted again that the emphasis here is on the description and illustration of the analysis and not on the numerical results, which are likely to change as additional expert assessments or new information become available. The outputs presented here are based on a limited number of expert elicitations at the end of 2014 in addition to estimates found in the open literature. All cases use 10,000 Monte Carlo simulation runs and assume no additional detections unless otherwise specified. The results are thus illustrative only, because other factors and additional information (e.g., other experts, intelligence information from different sources, etc.) was unavailable to us but could have been included.

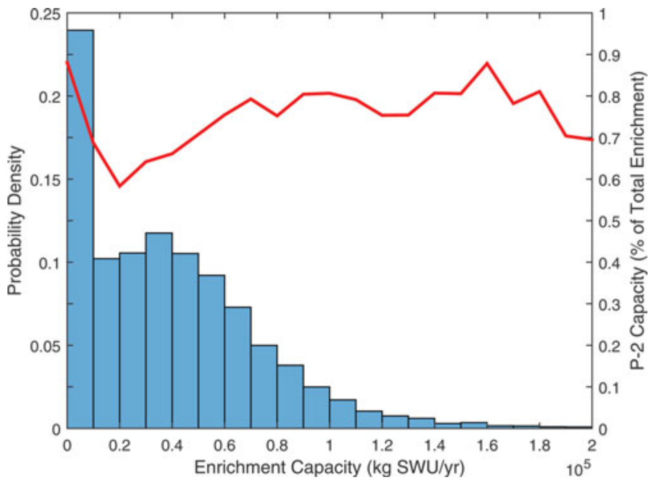

Figure 6: Illustrative enrichment capacity distribution (histogram) and fraction of total capacity from P-2 centrifuges (line).

Base Case Results

As illustrated in Figure 6, the resulting enrichment capacity distribution at the time of the study can be represented as a positively skewed probability density function (PDF) with a finite probability mass at zero and a main mode of approximately 35,000 kg SWU/yr.[24] This mode would translate into 4–6 bombs annually or enough enrichment capacity to support electricity production from nearly five reactors of the type of 25–30 MWe LWR reactors currently under construction.[25]

The same result (i.e., the illustrative enrichment capacity random variable in 2014) also can be represented by its cumulative distribution function (CDF). As shown in Figure 7, this CDF shows the probability that North Korea's uranium enrichment capacity at current centrifuge facilities (or those operational within one year) is less than or equal to the values shown on the x axis. This graph, based on expert assessment of disaggregated estimates, shows that the probability of a near-zero capacity of uranium enrichment is about 14 percent. Independently, Hecker assessed the probability of such a "Potemkin village" scenario (i.e., the centrifuges and control room shown to him were a mock-up) at less than five percent.[26] In the model, these zero-capacity scenarios occur when the North Koreans cannot produce or procure either ring magnets or special oils, which are both necessary to operate centrifuges. This graph also shows that there is an 81 percent chance that the total enrichment capacity is greater than 8,000 kg SWU/yr., which is the claimed output of the facility shown to Hecker.[27]

Examining which component stockpile constraint is binding for a particular scenario (i.e., corresponding to a Monte Carlo run) permits identification

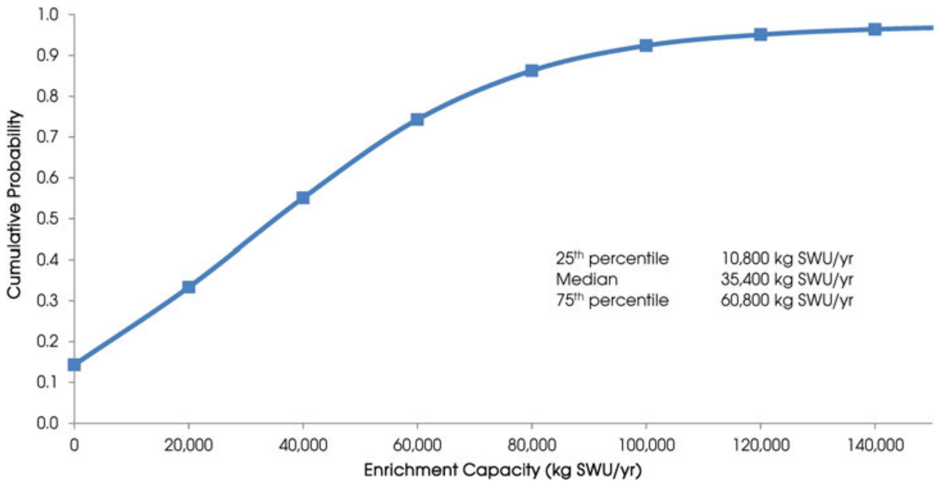

Figure 7: Cumulative distribution function for North Korea's uranium enrichment capacity, which plots the probability that the capacity is less than or equal to the x-axis value.

of the components that are bottlenecks for scaling up enriched uranium production. As stated earlier, these computations are based on current expert estimates and should be updated with new information as it becomes available. For a Monte Carlo simulation of 10,000 runs, the percentage of runs that a constraint is binding can be interpreted as the probability that each stockpile is a limiting factor for a marginal expansion of the program. Figure 8 suggests that the three most likely bottlenecks are pivot bearings (with a 0.36 probability of being the binding constraint), maraging steel, and high-strength aluminum. At the opposite extreme, controller units are very unlikely to be a limiting factor given the ease of obtaining these units and the relatively small number needed in centrifuge facilities.

Another result of interest is the number of centrifuges that North Korea can produce given their stockpiles of components and raw materials. The P-2 design, depending on its rotor speed and length, enriches many times the amount per centrifuge as the P-1. Therefore, the model assumes that North Korea would produce as many P-2s as possible given the available maraging steel. If aluminum is the binding constraint (i.e., if a sufficient supply of maraging steel is available), only P-2 centrifuges are produced. If maraging steel is the binding constraint, more P-1s are typically produced out of the excess aluminum. Optimal centrifuge production consists largely of enrichment from P-2s, as indicated in Figure 6. This result, however, depends on the probabilities of indigenous production displayed in Table A3. These values are highly uncertain and, hence, future work should systematically address the sensitivity of these results to these production capabilities and their correlations.

Sensitivity Analyses

Having identified potential stockpile bottlenecks, the next problem is to determine the sensitivity of the outputs to stockpile inventories. A stockpile sensitivity analysis was performed by removing the binding constraints one at a time, beginning with the most frequently binding, and recalculating the enrichment capacity distribution. These results display the sensitivity of the model but also show North Korea's potential for scaling up uranium enrichment if they can obtain additional components.

Figure 9 illustrates how the cumulative distribution function of North Korea's enrichment capacity changes as binding constraints are relaxed. After removing the constraint on pivot bearings (i.e., the most frequently binding constraint), the median of the capacity increases from 35,400 to 43,300 kg SWU/year.[28] After removing the constraints on all materials availability, the median increases to over 170,000 kg SWU/yr. The implication is that, if North Korea could obtain abundant maraging steel, high-strength aluminum, and the ability to manufacture pivot bearings, they could greatly scale up their uranium enrichment capacity, perhaps at another location.[29]

Dependent Exporter Probabilities

In the base case, it is assumed that the probability that any given country exported nuclear materials to North Korea was independent across materials. For this subsection, this assumption is relaxed for maraging steel and high-strength aluminum exports to reflect more accurately the experts' beliefs regarding the tendency of countries to ignore export restrictions if they be-

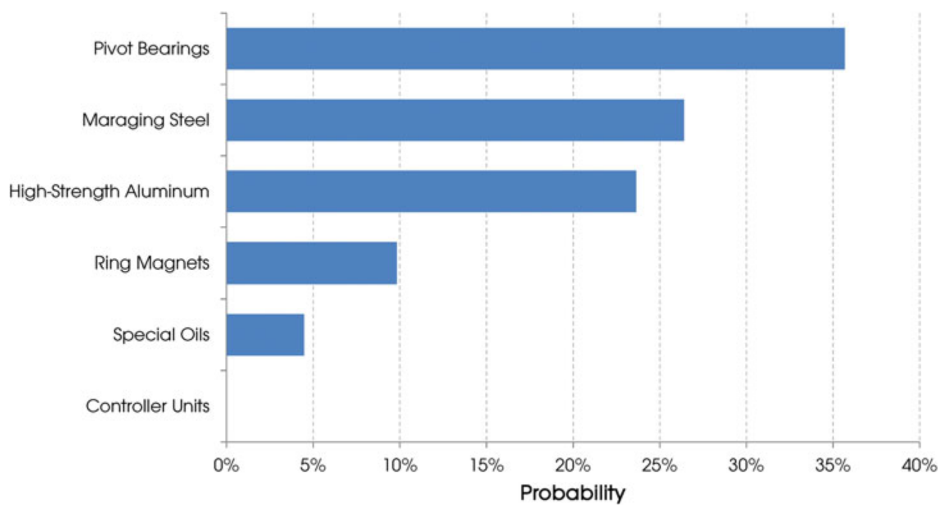

Figure 8: Probability that the availability of each stockpile is a binding constraint.

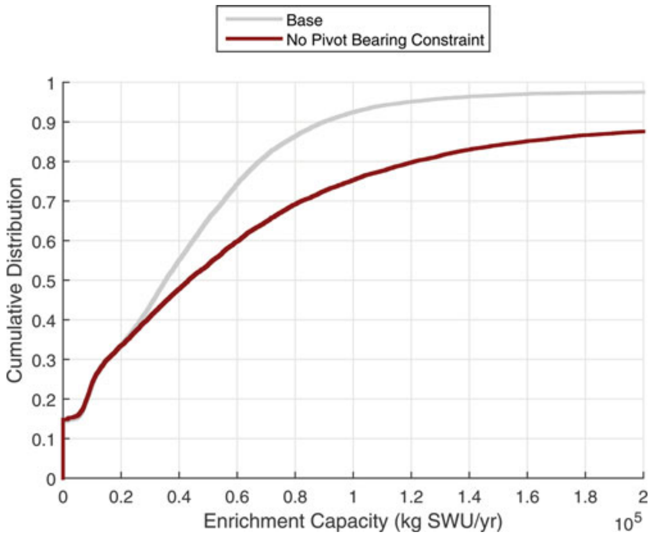

Figure 9: Cumulative distribution of the enrichment capacity with binding constraints removed.

lieve that others are also "cheating," and because earlier results in this section determined that North Korea's enrichment capacity is sensitive to stockpiles of these materials. Relaxing the exporter-independence assumption required that the experts assess the probability that country two exported nuclear materials to North Korea given that country one did or did not, that country three exported given that countries one and two did, that neither did, that two did but one did not, that one did but two did not, and so forth. The order of assessment does not matter as long as all possible combinations of the five potential exporter countries are assessed. For both materials, they chose to first assess the probability that Europe/Japan exported, then assess the role of Pakistan (given Europe/Japan did and did not export), then of Russia (given what the others could do), then of Iran, and finally of China.

The effect of considering dependences among potential nuclear exporters to North Korea is to increase both the mean and the variance of the probability of North Korea's enrichment capacity. While the base case the probability that enrichment capacity exceeds 8,000 kg SWU/year (the capacity claimed by North Korea) is 81 percent and the median enrichment capacity is 35,400 kg SWU/year, those same metrics for the case with dependent exporter probabilities are 86 percent and 52,200 kg SWU/year respectively, as shown in Figure 10. The intuition behind these shifts is that the experts believe that it is very likely that specific countries exported nuclear materials to North Korea in the past, and this occurrence, in turn, somewhat increased the likelihood that other countries also exported to North Korea (although not necessarily across the board).

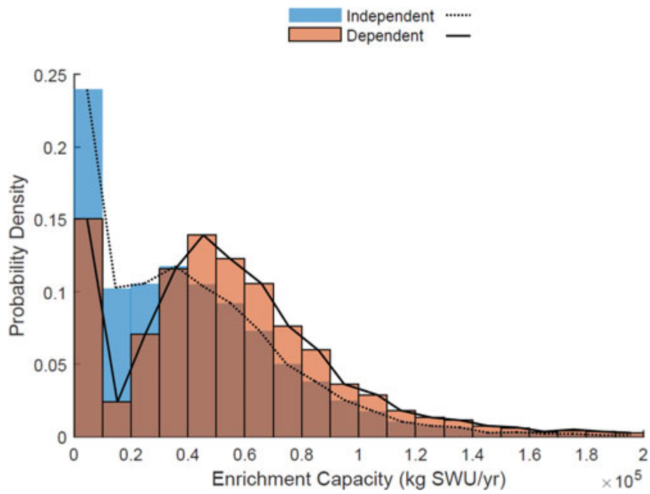

Figure 10: Comparison of probability density functions of North Korean uranium enrichment capacity with independent versus dependent exporters.

In other words, whereas exporter independence implied that North Korea was most likely able to enrich enough uranium for up to six bombs per year (mode of the distribution), relaxing this independence assumption increases North Korea's most likely output to up to nine bombs worth of uranium per year. More HEU allows North Korea to scale up its nuclear arsenal more rapidly and gives it more flexibility to conduct nuclear tests. North Korea will likely need to conduct more nuclear tests to miniaturize its devices in order to fit on top of one of their missiles. However, testing is likely more constrained by political forces than lack of fissile materials if the enrichment capacity is close to either estimate. The model presented here can be extended to include a range of dependencies among imports of critical elements of the whole North Korean nuclear program. This expansion depends on computer capabilities and ability to encode expert opinions but not on the model structure (Figure 11).

Scenario Analyses

The base model can be expanded to explore the effects of a diverse array of hypothetical scenarios. The motivation behind the scenario analyses in this section is twofold. First, they show the implications of altering the assistance that North Korea receives from international parties. Second, they demonstrate the flexibility of the model, and in particular, how the use of Bayes nets allows incorporation of new intelligence information related to shipping centrifuge components and materials (if one assumes that the likelihood functions

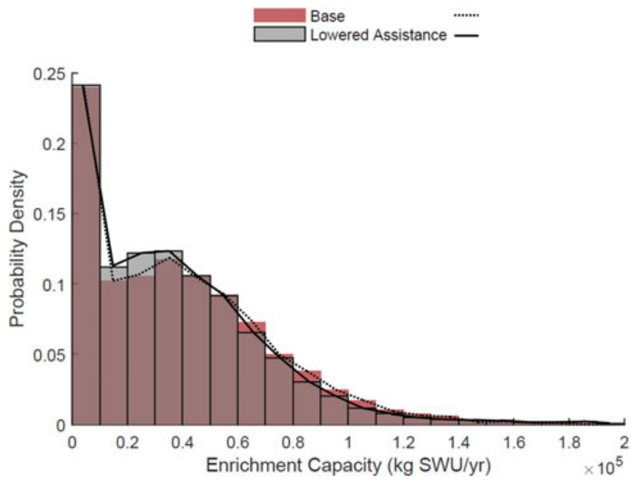

Figure 11: Comparison of probability density functions of North Korean uranium enrichment capacity between the base case and the reduced Chinese assistance scenario.

are constant). Here, three scenarios illustrate the ways in which the model can be adapted to reflect changes in policies or detections of shipments.

Assistance

The first scenario explores the effect of decreasing the probability of cooperation between North Korea and China for all components and materials. In this sensitivity analysis, the probability of Chinese assistance is decreased to 0.5 to determine its influence on the probability distribution of North Korea's ability to enrich uranium. Note that this probability of assistance has no direct implication for intergovernmental relationships but rather characterizes the chances that Chinese companies may serve as a point of origin for the shipment of components or materials.

Lower probabilities of Chinese assistance lead to lower probability mass in the right tail of the enrichment capacity distribution compared with the base case. However, the model and elicitations in Table A2 indicate that export parties in other countries may contribute centrifuge equipment and materials, even when the probability of Chinese assistance is lower, and in spite of the fact that recent United Nations Security Council sanctions have likely restricted imports from other countries. Like other results, these insights are linked with the elicited values for material and component acquisition by country of origin in Table A2, and updated intelligence assessments for these quantities may offer different qualitative conclusions.

The second assistance scenario illustrates what might happen to the North Korean centrifuge program if one component or material became much more difficult to obtain. In this example, the assistance probability for controller units is reduced to simulate placing these components on a dual-use list, since the lack of monitoring or export controls for this component contributed to the fact that it is never a binding constraint. To test the effects of this hypothesis, the probability of obtaining controller units from China, Europe, or Russia is reduced from 0.95 percent to 0.05, and from Iran or Pakistan from 0.95 to 0.50. The probabilities of assistance from these regions for all other components are kept unchanged. The changes that result in the probability density function of North Korea's uranium enrichment capacity are very small, which reflects the opinion of the experts that it would be easy to obtain controller units from a variety of sources. This analysis supports the view that although export controls are necessary, they typically do not stop a determined proliferator.

Detection

To illustrate the sub-model of component transfer detection, the hypothetical scenario considered here shows the change in the probability distribution of enrichment capacity if transportation of high-strength aluminum by sea originating from China is detected with certainty. Figure 12 shows

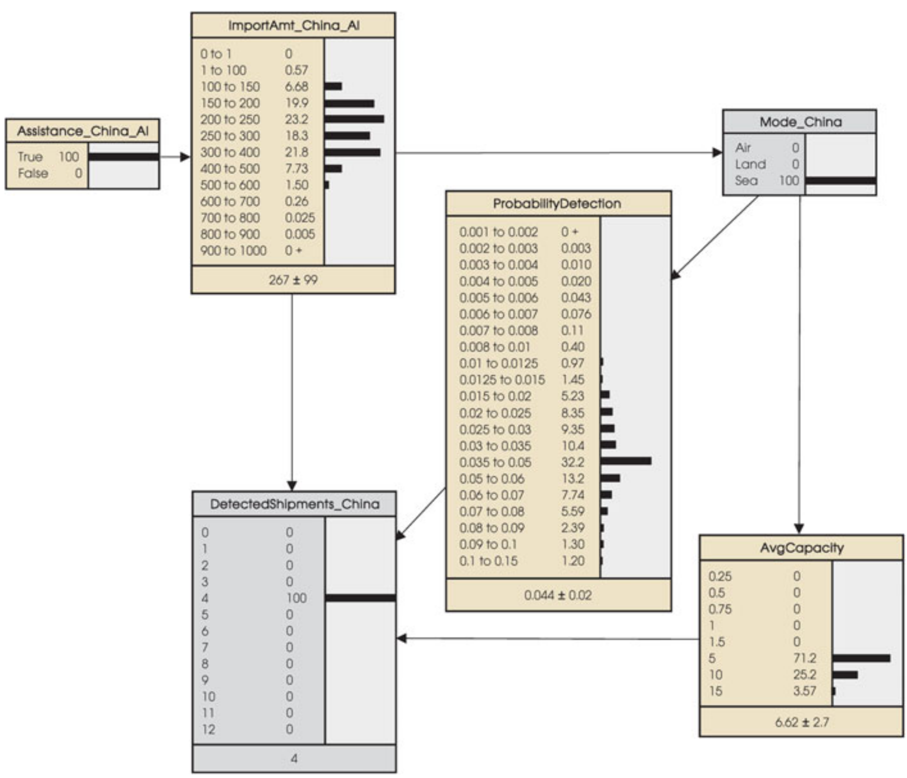

Figure 12: Updated Bayes net in Netica™ showing the sub-model for a detection scenario of sea shipments from China.

the detection sub-model for the amount of aluminum procured from Chinese sources. Before maritime detections, the prior distribution of the amount of aluminum originating in China shows a 0.58 probability that any shipment of any type of aluminum shipment goes undetected, a 0.20 probability of detecting one shipment, and a very low chance of detecting more than one shipment.

Assuming that four sea shipments of aluminum originating in China have been observed, Figure 12 shows the updated probability distributions using a Bayes net in Netica™. Note that, for all distributions other than that of the amount of aluminum originating in China, nothing is changed. However, the distribution of the amount of aluminum received from China has changed, with its expected value increasing by 400 tons. This new distribution can then be propagated through the optimization model to illustrate the effect of that updating on the distribution of North Korea's uranium enrichment capacity, as shown in Figure 13.

The distribution is displaced to the right, indicating that there are higher probabilities for greater enrichment capacities. This result makes sense in the context of the model given aluminum's frequency as a binding constraint. This scenario illustrates how shipment detections can have significant effects on the probability distribution of enrichment capacity. Direct or indirect detection is one of the primary ways for the international community to gather information on North Korea's centrifuge program. Therefore, the ability to use detections to update distributions of the amount of components or mate-

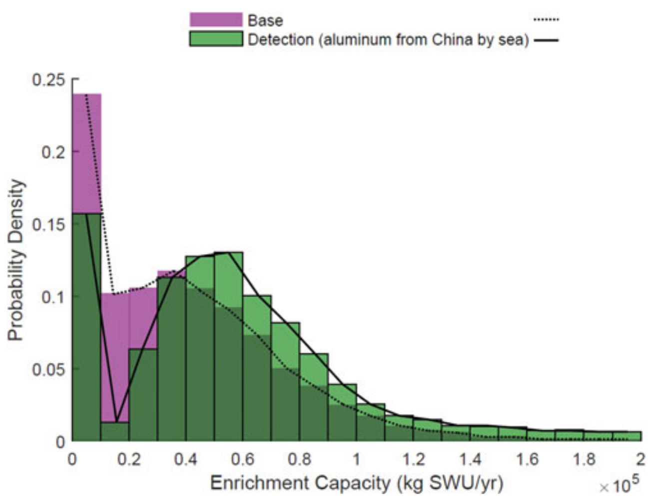

Figure 13: Comparison of North Korean uranium enrichment capacity probability density functions between the base case and the detection scenario (i.e., in which four sea shipments of aluminum from China are observed).

rials available to North Korea is particularly important for policy-makers and analysts.

CONCLUSIONS

The analysis detailed above offers a systematic approach to a probabilistic assessment of the size of North Korea's uranium enrichment program and identifies potential bottlenecks for scaling up production. The basic random inputs are the materials and components available to (but could be limiting factors in) centrifuge production. A deterministic optimization model is used to estimate the size of North Korea's uranium enrichment program as a function of uncertain component availability. Bayes nets are used to generate probability distributions of the sizes of stockpiled centrifuge components available for use to manufacture uranium centrifuges. A Monte Carlo simulation is used to propagate uncertainties about stockpiles of centrifuge components through the deterministic model and to generate the corresponding distribution of uranium enrichment capacity. All values presented here are illustrative and rely on limited expert opinions as of 2014. Very little is known about the North Korean nuclear program. Therefore, these figures can be considered as reasonable *open-source* estimates for North Korea's current capabilities. For decision support, however, these values should be updated by the appropriate organizations with the best available information, which may include classified estimates.

The current analysis, based on expert input about potential materials and component acquisition, supports the intuition of one of the authors (Hecker) that North Koreans have greater enrichment capacity than the 8,000 kg SWU that they claimed to have in 2010. The analysis indicates that the probability that the facility at Yongbyon has an enrichment capacity greater than or equal to 8,000 kg SWU is about 80 percent. The probability distributions for that capacity show that one cannot ignore the possibility that North Korea has an annual enrichment capacity of several tens of thousands kg SWU. The results also demonstrate that one cannot discount the fact that North Korea may have zero (or close to zero) capability, although recent activity at the fuel fabrication facility (and more recent data not included here) makes this even less likely than the illustrative results of the model suggest.

The three most likely bottlenecks to higher production capacity are the availability of pivot bearings (with a 0.36 probability of being the binding constraint), high-strength aluminum, and maraging steel. The stockpile sensitivity analysis demonstrates that, if North Koreans could obtain abundant high-strength aluminum or maraging steel and the ability to manufacture pivot bearings, they could greatly scale up operations. Abundant maraging steel would permit a significant increase in enrichment capacity through the P-2 route, but if unlimited quantities of high-strength aluminum were available instead, extensive production of P-1s could also lead to high enrichment capacities. Although the exact magnitude of this increase depends on the input distributions, this analysis can be performed easily using the model. Future work could also involve the case in which North Korea has developed the capability of making 250-grade maraging steel indigenously (at a probability > 0.2 shown in Table A3) and is willing to settle for lower rotor speeds. These features could then be included in the model by modifying the probabilities displayed in Tables A2 and A3, and the results would change accordingly. The detection sub-model offers the ability to include new intelligence information about the transfer of components to North Korea and to update the probability distributions without the need for reassessment of the signals (assumed here to be of constant perfect quality). A hypothetical scenario examined here shows changes in the distribution of enrichment capacity if transportation by sea of high-strength aluminum originating from China is detected. In future work, one may want to incorporate a more sophisticated Bayesian model to account for the quality of detection signals, thus relaxing the assumption of perfect observations.

The model presented here can thus be most useful to strategy and policymakers in support of decisions involving North Korea's ability to produce nuclear weapons. It will, however, need constant, real-time maintenance and updating to reflect the best available information at the time of those decisions.

NOTES AND REFERENCES

1. D. Albright, and P. Brannan, "Taking Stock: North Korea's Uranium Enrichment Program," Institute for Science and International Security, 8 October 2010, http://isis-online.org/uploads/isis-reports/documents/ISIS_DPRK_UEP.pdf.
2. S. Hecker, and R. Carlin, "North Korea in 2011: Countdown to Kim il-Sung's Centenary," *The Bulletin of Atomic Scientists* (January/February 2012).
3. D. Albright, and P. Brannan, "Taking Stock: North Korea's Uranium Enrichment Program."
4. S. Hecker, "A Return Trip to North Korea's Yongbyon Nuclear Complex," Center for International Security and Cooperation, Stanford University, November 2010, http://iis-db.stanford.edu/pubs/23035/HeckerYongbyon.pdf.
5. S. Hecker, "Interview with Siegfried Hecker: North Korea Complicates the Long-Term Picture," *The Bulletin of the Atomic Scientists* 5 (2013), http://www.thebulletin.org/web-edition/features/interview-siegfried-hecker-north-korea-complicates-the-long-term-picture.
6. A Bayes net (Bayesian network, or influence diagram) is a compact representation of a classic event tree, where uncertainties are represented by the marginal probability distributions of some random variables and by the conditional distributions of dependent ones. Like its homomorphic event tree, it allows computation of the distribution of an outcome variable.
7. The Monte Carlo simulation draws random values from the distributions of each input (i.e., stockpiles of the different centrifuge components). Each of these sets of

random values defines a scenario and becomes an input to the deterministic centrifuge optimization model. For each scenario, this model computes the corresponding maximum enrichment capacity (in kg SWU/yr.) that can be obtained given the availability of the different components. The sampling and optimization are repeated to construct a probability distribution of the North Korean enrichment capacity at a given time based on available information.

8. D. Albright and M. Hibbs, "Iraq's Shop-Till-You Drop Nuclear Program," *The Bulletin of the Atomic Scientists* (1992).
9. Materials such as maraging steel and aluminum have to be processed after they are imported. North Korea's capability to accomplish that process is uncertain.
10. Maraging steel is a low-carbon, ultra-high strength iron alloy that includes about 18 percent nickel as well as cobalt, molybdenum, titanium, and other minor alloying elements. It has very good fracture toughness, is readily weldable, and easily heat treated to very high strength. The 250 and 350 grades refer to the yield strengths in thousands of pounds per square inch.
11. The availability of strontium or barium ferrite magnets in countries like China may also be very high. Iran placed an order in 2011 for 100,000 such magnets for their P-1 centrifuges with a Chinese company (although it is not known whether it was ever filled) according to David Albright, "Ring Magnets for IR-1 Centrifuges," Institute for Science and International Security, 13 February 2013, http://isis-online.org/uploads/isis-reports/documents/iran_ring_magnet_13Feb2013.pdf.
12. M. Granger Morgan, and M. Henrion, *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis* (Cambridge: Cambridge University Press, 1990).
13. D. Albright, and C. Walrond, "North Korea's Estimated Stocks of Plutonium and Weapon-Grade Uranium," Institute for Science and International Security, 16 August 2012, http://isis-online.org/uploads/isis-reports/documents/dprk_fissile_material_production_16Aug2012.pdf.
14. Note that these illustrative assessments of enrichment capacity are not necessarily superior, since the model in this paper requires many parameter estimates with a limited amount of open-source information and relies on a few experts.
15. The shadow price of each constraint is defined as the increase in enrichment capacity that would result from relaxing by one unit the binding constraint for each element.
16. R. S. Kemp, "Gas Centrifuge Theory and Development: A Review of U.S. Programs," *Science and Global Security* 17 (2009): 1–19.
17. The experimental efficiency aggregates inefficiencies not quantified elsewhere in the equation such as disturbances to the countercurrent flow and turbulent mixing. This efficiency is unique to each centrifuge design and improves with centrifuge production experience.
18. S. Hecker, "A Return Trip to North Korea's Yongbyon Nuclear Complex."
19. R. S. Kemp, "Gas Centrifuge Theory and Development," 1–19.
20. *Technical knowledge* is a binary, random variable that represents whether or not North Korea has the technical know-how to construct a specific component.
21. *Manufacturing capacity* is a binary, random variable that represents whether or not North Korea has the necessary equipment to produce a specific component.
22. *Assistance* is a binary, random variable that represents whether or not a specific country or region has served as the origin of a type of component or material and does not necessarily indicate an intergovernmental relationship between North Korea and the country of interest.

23. The gamma distribution was chosen to fit the expert's assessments because its general shape has features that reflect the expert's beliefs. It can have no density at zero, packs most of its density toward lower quantities, but includes a large tail stretching toward infinity. Examples of fitting gamma distributions to expert assessments can be found in the Appendix.
24. In this analysis, the "mode" is typically used as a measure of the central tendency of enrichment capacity distributions. The mean (average) is a less salient statistic in this context due to the large influence of low-probability, high-capacity samples on its value. Nonetheless, the tail of the distribution, which heavily influences the mean, is important because it represents low-probability, high-consequence situations that may have to be considered in policy decisions. The mode is the value that corresponds to the maximum of the probability distribution. It represents the enrichment capacity value that is most often sampled from the discrete distribution shown in Figure 6.
25. These approximations are based on the uranium enrichment calculator found at http://www.wise-uranium.org/nfcue.html.
26. Hecker has revised his estimate to less than one or two percent because of the significant amount of construction that is now visible from overhead imagery at the fuel fabrication facility, which he believes to support the enrichment program. In addition, exterior construction of the experimental LWR appeared complete in late 2013 and will likely go operational some time in 2015. It requires LEU fuel, which likely has to be produced domestically. Hecker's estimates demonstrate the importance of on-site visits and observations since the North Korean nuclear program has historically been underestimated based on indirect information and estimates.
27. S. Hecker, "A Return Trip to North Korea's Yongbyon Nuclear Complex."
28. Removing this constraint is equivalent to assuming that North Korea obtains the indigenous capacity for producing pivot bearings with probability one.
29. Although the exact magnitude of this increase is dependent on the input distributions, this analysis can be performed quite easily in this modeling framework.
30. Matthias Englert, personal interview on 18 April 2011. R. Scott Kemp, personal interview on 26 September 2013.
31. D. Albright and C. Walrond, "Iran's Gas Centrifuge Program: Taking Stock," Institute for Science and International Security, 11 February 2010, http://isis-online.org/uploads/isis-reports/documents/Natanz_Operation_11Feb2010.pdf; D. Albright, and P. Brannan, "Taking Stock: North Korea's Uranium Enrichment Program."
32. *Amount received by mode of transportation* is a continuous, random variable that represents the total amount of a component across all shipments to North Korea originating from a country and sent by way of air, sea, or land.
33. *Probability of detection* is a continuous, random variable that represents the probability of detecting any given shipment of a material or component.
34. *Average amount per shipment* is a discrete, random variable that represents the average (mean) amount of components or materials aboard a shipment to North Korea.
35. *Number of shipments detected* is a discrete, random variable that follows a binomial probability distribution once the amount received, probability of detection, and average amount per shipment have been determined.
36. The binomial coefficient reflects the fact that, for a total of six shipments that are categorized as either detected or undetected, there are more ways that three detected and three undetected can occur than all detected or all undetected. This

is analogous to craps, where a person rolling the dice is more likely to roll a seven than snake eyes.

37. The alternative approach to rounding would be to add an additional probabilistic dependence of $\tilde{q}$ on $\tilde{a}$, and then to place non-zero probability of $\tilde{q}|\tilde{a} = a$ only on integer multiples of a. This method, however, would require significantly more expert assessments and is therefore more difficult to implement.
38. This assumption implies, for example, that the probability of detection does not change under repeated interdictions (even if the smuggler adjusted routes in pursuit of unmonitored channels) and that the capabilities of intelligence agencies do not change, assumptions that can be relaxed by allowing variations of the detection model parameters.
39. E. Paté-Cornell, "Fusion of Intelligence Information: A Bayesian Approach," *Risk Analysis* 22 (2002): 445–54; Gary McClelland, "Use of Signal Detection Theory as a Tool for Enhancing Performance and Evaluating Tradecraft in Intelligence Analysis," in *Intelligence Analysis: Behavioral and Social Scientific Foundations*, ed. Baruch Fischhoff and Cherie Chauvin (Washington: National Academies Press, 2011), 83–99.

APPENDIX

Model Data

The model required extensive data, which were gathered from a variety of sources. The Bayes net models needed expert elicitations of probability distributions of component and material quantities by country of origin. Centrifuge construction data for the optimization model also rely heavily both on expert elicitations,[30] as well as estimates from published articles.[31] Table A1 lists a few values for the quantity of processed components and materials (as opposed to their raw forms) required for centrifuges P-1 and P-2.

It is assumed that high-strength aluminum alloys are only used for the rotors of the P-1 design. The casings of the P-1s can be constructed of lower-grade alloys, which are assumed to be available to the North Koreans (Figure A1).

Mathematical Formulation of Centrifuge Optimization Model

This analysis uses a centrifuge optimization model to determine the optimal number of centrifuges of each type that can be constructed given available stockpiles of materials (a result that comes from sampling the Bayes net distributions). The objective of this optimization is to maximize the total uranium enrichment capacity (in kg SWU/yr.) of North Korea's program

Table A1: Quantities of components and materials required per centrifuges of different designs

Component/material	Quantity required (per centrifuge)	
	P-1	*P-2*
Pivot bearings	1	1
Controller units	1 / 1,000	1 / 1,000
High-strength aluminum	0.9 kg	0 kg
Maraging steel	0.75 kg	6 kg
Special oils	1	1

One S7-417 PLC from Siemens is estimated to control around six cascades with 164 centrifuges each.

formulated as:

$$\begin{aligned} \max f(\mathbf{x}) &= \mathbf{c}^T\mathbf{x} \\ \text{s.t. } \mathrm{P}\boldsymbol{x} &\le \mathbf{s} \\ \mathbf{x} &\ge 0 \end{aligned} \qquad \text{(A1)}$$

Where:

c vector of potential uranium enrichment capacities (in kg SWU/yr.) for different centrifuge designs;

x vector of potential number of centrifuges of different designs produced from available components (i.e., vector of design variables);

P matrix of parameter values of required components and materials required for different centrifuge designs;

s vector of available stockpiles for centrifuge components and materials.

The model f(x) is additive and simply means that the enrichment capacity is the sum of the quantities of enriched uranium that can be produced by all available centrifuges. The constraints mean that the total amount of components used to manufacture the centrifuges is less than or equal to the quantities available. This model also implies that the vector c has been optimized so that the number of centrifuges of different types provides the maximum enrichment capacity.

Elicitation Data

The data used for illustration in this article represent the expert opinions of Siegfried Hecker and Chaim Braun, Center for International Security and Cooperation, Stanford University (Figure A2).

The amounts cited in the table are estimates for acquired raw materials, in contrast to the amounts cited in Table A1, which are for finished products. The judgment of experts was used to estimate processing losses. Note that 350-grade maraging steel is on the Nuclear Suppliers Group dual-use export control list, but 250 grade is not because it has many other industrial applications and therefore may have been somewhat easier to acquire. Only a few countries have mastered the ability to produce the 350 grade. The 250 grade requires less stringent melting and vacuum controls as well as fewer controls on the thermo-mechanical processing. Therefore, it is both easier to import and

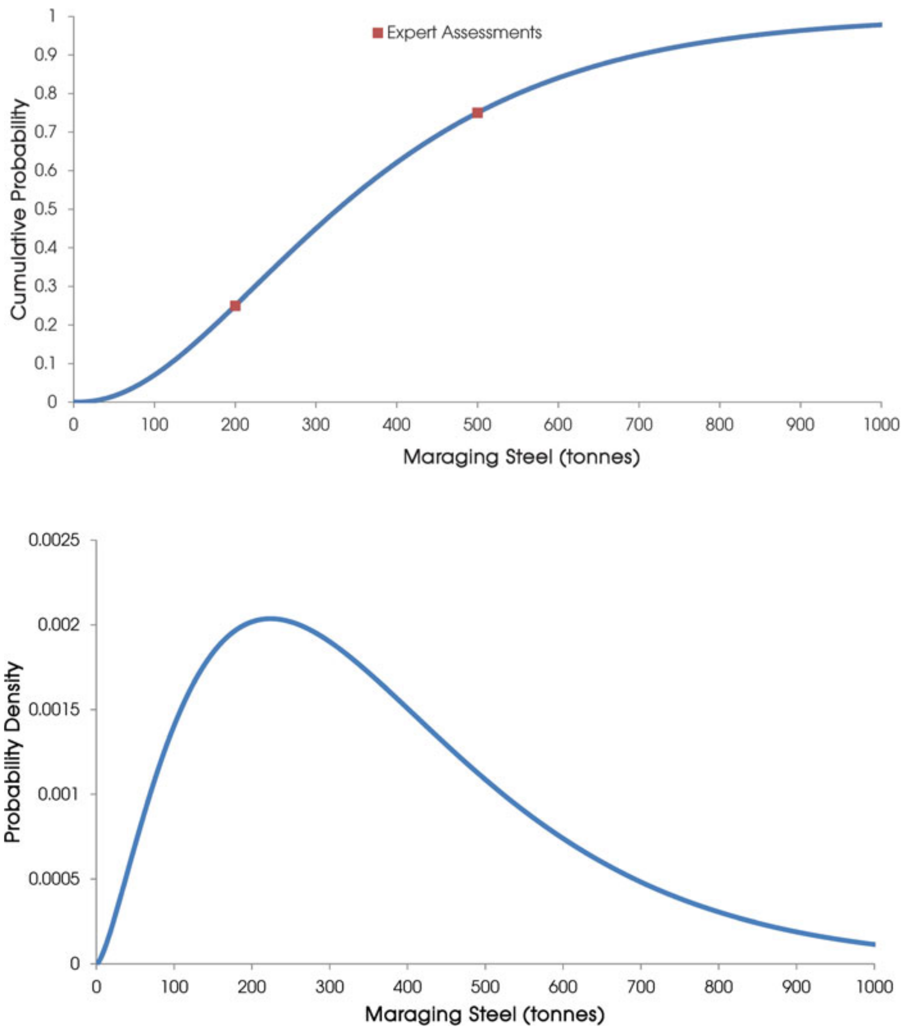

Figure A1: Gamma distribution fitting to expert assessments for maraging steel imports from Russia cumulative distribution function (top) and probability density function (bottom).

		ImportAmt | Assistance_Country = T			x1 < ImportAmt_Country < x2		Probability of Detection			Amt Carried by Mode		
Country	P(Assistance)	Distribution	α	β	a	b	p_air	p_land	p_sea	x_air	x_land	x_sea
Russia	0.85	Gamma	2.46	153.44	14000	35000	0.001	0.005	0.02	1000	500	10000
China	0.95	Gamma	0.97	37.37	15000	40000	0.001	0.005	0.02	1000	500	10000
Europe	0.8	Gamma	1.79	123.59	5000	15000	0.001	0.005	0.02	1000	500	10000
Pakistan	0.95	Gamma	0.97	37.37	500	4000	0.001	0.005	0.02	1000	500	10000
Iran	0.9	Gamma			500	4000	0.001	0.005	0.02	1000	500	10000

MODE OF TRANSPORTATION ASSESSMENTS

Russia	P(Air)	P(Land)	P(Sea)	SUM
ImportAmt_Country < a	0.05	0.6	0.35	1
a < ImportAmt_Country < b	0.01	0.5	0.49	1
ImportAmt_Country > b	0.01	0.25	0.74	1

China	P(Air)	P(Land)	P(Sea)	
ImportAmt_Country < a	0.05	0.6	0.35	1
a < ImportAmt_Country < b	0.01	0.5	0.49	1
ImportAmt_Country > b	0.01	0.25	0.74	1

Europe	P(Air)	P(Land)	P(Sea)	
ImportAmt_Country < a	0.05	0.6	0.35	1
a < ImportAmt_Country < b	0.01	0.5	0.49	1
ImportAmt_Country > b	0.01	0.25	0.74	1

Pakistan	P(Air)	P(Land)	P(Sea)	
ImportAmt_Country < a	0.05	0.6	0.35	1
a < ImportAmt_Country < b	0.01	0.5	0.49	1
ImportAmt_Country > b	0.01	0.25	0.74	1

Iran	P(Air)	P(Land)	P(Sea)	
ImportAmt_Country < a	0.05	0.6	0.35	1
a < ImportAmt_Country < b	0.01	0.5	0.49	1
ImportAmt_Country > b	0.01	0.25	0.74	1

INDIGENOUS CAPABILITY ASSESSMENTS

	P(True)	P(False)
Technical Knowledge	0.4	0.6

Manufacturing	P(Manufacturing |TK = T)	P(Manufacturing| TK = F)
TRUE	0.1	0.01
FALSE	0.9	0.99

P(IndigCap)	0.04

Figure A2: Probabilities assessed by experts as model inputs, and parameters of the gamma distributions fitted accordingly.

North Korea may have developed indigenous capability to produce 250 grade. There is little information that 250 grade has been used in commercial centrifuge cascades, but for small production operations (e.g., those required to produce HEU for a few weapons), the approximate 20 to 30 percent reduction in throughput may be considered tolerable (Table A3).

Table A2: Probability assessments for material and component acquisition by country of origin

Material	Percentile	Russia	China	Europe/Japan	Iran	Pakistan
7075 Aluminum (metric tons)						
	10	50	3	25	1	3
	90	150	75	125	13	25
250-grade maraging steel (metric tons)						
	10	25	3	3	1	3
	90	125	50	75	3	13
350-grade maraging steel (metric tons)						
	10	25	1	3	1	1
	90	125	13	50	1	8
Pivot bearings (count)						
	10	1,000	2,000	5,000	100	100
	90	10,000	10,000	10,000	300	1,000
Controlling units (count)						
	10	2,000	2,000	5,000	100	100
	90	10,000	10,000	10,000	300	300

A number of countries use 300-grade maraging steel for bellows. Its fabrication is also quite challenging; hence, for the purpose of our estimates, it is combined with 350 grade in this table.

Table A3: Probability of domestic production for centrifuge materials and components (expert opinions)

Material	Probability
7075 aluminum	0.1
250-grade maraging steel	0.2
350-grade maraging steel	0.05
Special oils	0.95
Pivot bearings	0.2
Controlling units	0.2
Ring magnets	0.9

For the model's representation of special oils and ring magnets, any private contractual arrangements with foreign companies to make ring magnets from Chinese samarium and also acquisition of special oils from various sources are characterized as "domestic production."

Detection Sub-model Formulation

The following notation is used to define the detection sub-model.

- Discrete random variables are denoted as capitalized letters under a tilde
- Continuous random variables are denoted as lowercase letters under a tilde
- Conditional probability distribution of random variable $\tilde{X}$ given that random variable $\tilde{Y}$ is equal to value y is denoted as $P(\tilde{X}|\tilde{Y}=y)$
- Coefficient of the b^{th} term in the algebraic expansion of $(x+y)^n$ is denoted $\binom{n}{b}$

In addition to the *assistance* and *amount received* (i.e., uncertainties that appear throughout the probabilistic model), the detection sub-model includes the uncertainties: *the amount received by mode of transportation* ($\tilde{q}$)[32] *probability of detection* ($\tilde{p}$),[33] *average amount per shipment* ($\tilde{a}$),[34] and *number of shipments detected* ($\tilde{X}$).[35] The amounts received by each mode of transportation are summed together to get the total quantity received from a particular source.

The inputs required for that model are detection probabilities. The model currently provides inputs conditional on different modes of transit (including land, sea, and air shipments). For example, given a particular amount of high-strength aluminum imported to North Korea from China, the detection sub-model utilizes an assessment of the amount of material coming from each transportation mode. However, the model can easily accommodate mode-independent probabilities. Whether assessing mode-dependent probabilities will materially affect the results of analysis will depend in part on how countries are detecting smuggled shipments (e.g., if detection events occur from purchase order reporting or electronic communications instead of by customs agents). Small differences between the detection probabilities corresponding to two different modes of transportation can imply very different probability distributions over quantities of goods smuggled because of the different tonnage and volume available via each mode of transportation.

In the probabilistic model, the probability of detection is represented by three beta distributions, one for each mode of transportation. Currently, for each component, the probability of detection is thus represented by three beta distributions, one for each mode of transportation. These distributions have similar modes, but their variances increase as the mode of transportation shifts from air to land to sea, reflecting the belief that there is a greater probability of detection at sea than in the air. Given the mode of transportation, the detection sub-model also requires an assessment (either probabilistic or deterministic) of the average amount per shipment. Once all the required probability distributions have been encoded, the detection sub-model then assumes that the number of shipments detected ($\tilde{X}$) is binomially distributed given the amount received ($\tilde{q}$), the probability of detection ($\tilde{p}$) and the average amount per shipment ($\tilde{a}$) according to Equation 2.[36] The term $\lceil q/a \rceil$in Equation 2 is rounded upward to the next highest integer value.[37]

$$P(\tilde{X} = i | \tilde{Q} = q, \tilde{a} = a, \tilde{p} = p) = \binom{\lceil q/a \rceil}{i} p^i (1-p)^{\lceil \frac{q}{a} \rceil - i} \tag{A2}$$

Netica™ then permits computation of the probability distribution of the number of shipments detected (the distribution of the outcome of the Bayes net), which in turn permits inference of the probability distribution of the amount of a material or component that North Korea has received from foreign suppliers given a specific number of detection events. Because the likelihoods of signals given detection have already been encoded, this updating allows the inclusion of new information in the probabilistic model without additional elicitation of expert opinion after each detection.

In the probabilistic model described above, it is assumed that each detection event is an independent Bernoulli trial, which means that each event is independent of all others conditional on the realizations of random variables $\tilde{q}$, $\tilde{p}$, and $\tilde{a}$ and that the probability of detection and the average amount per shipment do not change over time.[38] In addition, the detection sub-model assumes that the probability of detection does not depend on the average amounts per shipment. Finally, it assumes that detections are perfect, meaning that observations from intelligence sources are assumed to be truthful and accurate. Methods to relax this final assumption are the subject of several papers in the literature.[39] One challenge in adopting these methods is that treating multiple imperfect detection signals as probabilistically dependent would make it too complex for an expert to assess the distributions of dependent detection variables for all materials or components. The model thus treats detections as perfect information and places the burden of evaluating detections on the analyst. He or she must judge *a priori* the veracity of each detection and include it in the detection count if estimated true (or exclude it otherwise).