

Analyzing a bioterror attack on the food supply: The case of botulinum toxin in milk

Lawrence M. Wein*[†] and Yifan Liu[‡]

*Graduate School of Business and [†]Institute for Computational and Mathematical Engineering, Stanford University, Stanford, CA 94305

Edited by Barry R. Bloom, Harvard University, Boston, MA, and approved April 20, 2005 (received for review November 16, 2004)

We developed a mathematical model of a cows-to-consumers supply chain associated with a single milk-processing facility that is the victim of a deliberate release of botulinum toxin. Because centralized storage and processing lead to substantial dilution of the toxin, a minimum amount of toxin is required for the release to do damage. Irreducible uncertainties regarding the dose–response curve prevent us from quantifying the minimum effective release. However, if terrorists can obtain enough toxin, and this may well be possible, then rapid distribution and consumption result in several hundred thousand poisoned individuals if detection from early symptomatics is not timely. Timely and specific in-process testing has the potential to eliminate the threat of this scenario at a cost of <1 cent per gallon and should be pursued aggressively. Investigation of improving the toxin inactivation rate of heat pasteurization without sacrificing taste or nutrition is warranted.

bioterrorism | mathematical modeling

Among bioterror attacks not involving genetic engineering, the three scenarios that arguably pose the greatest threats to humans are a smallpox attack, an airborne anthrax attack, and a release of botulinum toxin in cold drinks (1). The methods of dissemination in these three scenarios are, respectively, the person-to-person spread of a contagious disease, the outdoor dispersal of a highly durable and lethal agent, and the large-scale storage and production and rapid widespread distribution and consumption of beverages containing the most poisonous substance known. The first two scenarios have been the subject of recent systems modeling studies (2–5), and here we present a detailed systems analysis of the third scenario. For concreteness, we consider a release in the milk supply, which, in addition to its symbolic value as a target, is characterized by the rapid distribution of 20 billion gallons per year in the U.S.; indeed, two natural *Salmonella* outbreaks in the dairy industry each infected ≈200,000 people (6). Nonetheless, our methods are applicable to similar food products, such as fruit and vegetable juices, canned foods (e.g., processed tomato products), and perhaps grain-based and other foods possessing the bow-tie-shaped supply chain pictured in Fig. 1.

The Model

The mathematical model considers the flow of milk through a nine-stage cows-to-consumers supply chain associated with a single milk-processing facility (Fig. 1). *Supporting Appendix*, which is published as supporting information on the PNAS web site, contains a detailed mathematical formulation of the model, a discussion of the modeling assumptions, and the specification of parameter values, some of which are listed in Table 1. The supply-chain parameter values are representative of the California dairy industry, which produces >20% of the nation's milk (California dairy facts, www.dairyforum.org/cdf.html, accessed on May 18, 2004). In our model, cows are milked twice daily, and the milk from each farm is picked up once per day by a 5,500-gallon truck, which makes two round trips daily between various farms and the processing plant. Upon a truck's arrival at the processing plant, the milk is piped into one of several raw

milk silos, each capable of holding ≈50,000 gallons. Raw milk is piped into the processing facility, goes through a sequence of processes (e.g., separation, pasteurization, homogenization, and vitamin fortification), where each processing line may simultaneously receive milk from several silos, and is held in 10,000-gallon postpasteurization tanks before being bottled. In our base case, we assume that milk from different silos does not mix during downstream processing and relax this assumption later; although downstream mixing is physically possible at many facilities, it is not always done. Bottled milk is stored as finished-goods inventory before traveling through the downstream distribution channel, eventually being purchased and consumed.

We assume that botulinum toxin is deliberately released in either a holding tank at a dairy farm, a tanker truck transporting milk from a farm to the processing plant, or a raw milk silo at the processing facility. Each of these release locations leads to identical consequences, because the toxin is eventually well mixed throughout the contents of a raw milk silo. The crux of our analysis is to calculate the amount and toxin concentration of contaminated milk (see Fig. 4, which is published as supporting information on the PNAS web site). By California state law, a raw milk silo must be cleaned after 72 h of operation. During these 72 h, the silo is initially filled up, then replenished (i.e., simultaneously filled and drained) for most of the 72-h period, and finally drained empty by 72 h. Because the toxin concentration in the silo drops exponentially during the replenishment interval, each postpasteurization holding tank has a different concentration level. Moreover, the amount of contaminated milk and the concentration distribution are themselves random quantities, depending upon when in the 72-h silo operation cycle the deliberate release occurs. Because of the difficulty of a terrorist in scheduling the release for maximum impact, we assume the release occurs randomly throughout the filling and replenishment intervals and report the mean number of poisoned people averaged over the random release time within the cycle. Using heat-inactivation data for foods with similar pH (7), we estimate that the heat-pasteurization process [170°F (77°C) for 15 min] inactivates 68.4% of the toxin.

Each gallon of purchased milk is continuously consumed by four people (one child and three adults) over a 3.5-day period. Children aged 2–11 and adults have differential milk consumption rates and dose–response curves in our model. A probit dose–response model dictates the precise timing of each poisoning. Our dose–response relationship is based on scant human data ($ID_{50} = 1 \mu\text{g}$ for adults, $ID_{50} = 0.43 \mu\text{g}$ for children) (8, 9). The attack can be detected via either early symptomatics or in-process testing results, whichever occurs first. We assume the outbreak is detected when the 100th person develops symptoms [the incubation period, which is the interval between the time of poisoning and the onset of symptoms, is log normal with a median of 48 h and a dispersal factor of 1.5 (10)], and an

This paper was submitted directly (Track II) to the PNAS office.

Freely available online through the PNAS open access option.

[†]To whom correspondence should be addressed. E-mail: lwein@stanford.edu.

© 2005 by The National Academy of Sciences of the USA

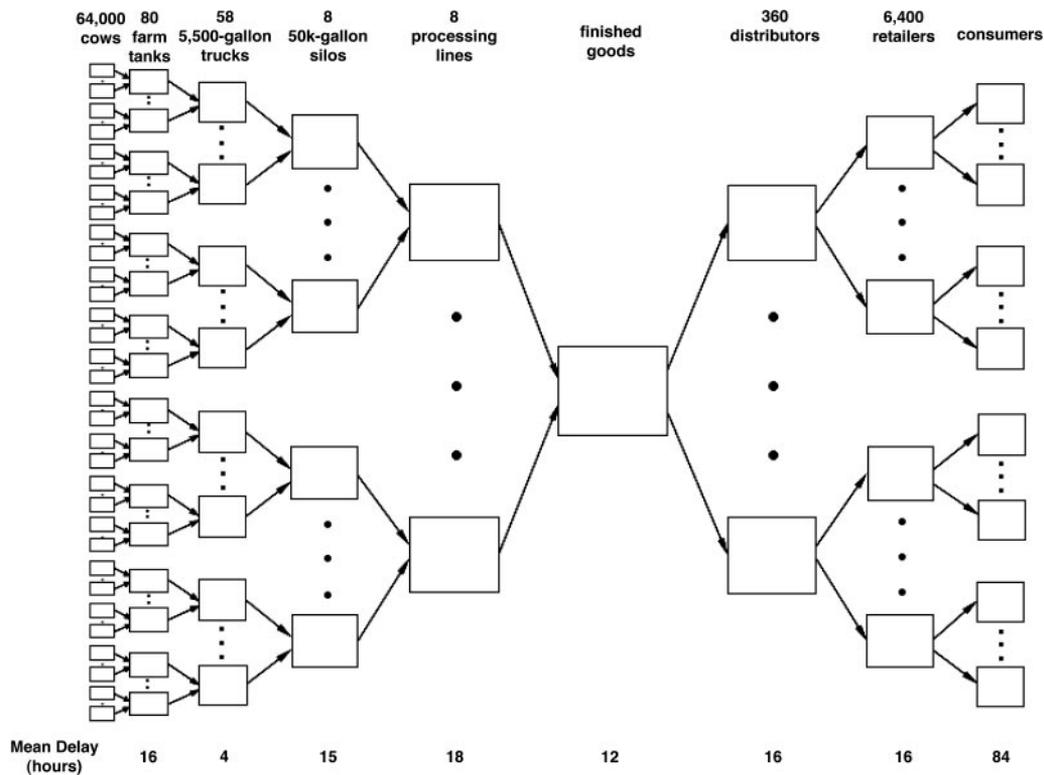


Fig. 1. The milk supply chain.

additional 24 h are required to identify the attack as being milkborne, at which time all consumption is halted. As with current antibiotic residue testing, we assume in-process botulinum testing is performed on milk from each truck just before the milk is piped into a raw milk silo at the processing facility. We have two tests at our disposal: the Food and Drug Administration-approved mouse assay with a detection limit of 16 pg/ml (11) and a testing delay of 48 h, and an ELISA test with a detection limit of 80 pg/ml (12) and a testing delay of 3 h. Because the mouse assay is not practical for widespread use (assays are processed at only several U.S. laboratories, and the mouse supply is limited), we assess two strategies: the ELISA test used in isolation (i.e., consumption is stopped after a positive ELISA result) and a sequential strategy in which the mouse assay is used as a confirmatory test after a positive ELISA result (i.e., consumption is halted after a positive mouse result). The latter strategy has a detection limit of 80 pg/ml and a testing delay of 51 h. The ELISA test in isolation is practical only if the test has an extremely small false-positive rate (no data have been published on ELISA test specificity in milk); otherwise, the sequential strategy is the only viable alternative.

Results

In the absence of any detection (i.e., every gallon of contaminated milk is consumed), the mean number of people who consume contaminated milk is 568,000 (Fig. 2). Less than 1 g of toxin is required to cause 100,000 mean casualties (i.e., poisoned individuals), and 10 g poison the great majority of the 568,000 consumers (Fig. 2). Most of the casualties occur on days 3–6, although they happen somewhat faster for larger releases, because less consumption is required for poisoning. Due to children's higher consumption rate and greater toxin sensitivity, the percentage of casualties who are children in Fig. 2 decreases from 99.97% for a 0.1-g release, to 61% for a 1-g release, to 28% for a 10-g release.

Early symptomatic detection avoids $\approx 2/3$ of the casualties in Fig. 2 (see Fig. 3) but still allows $>100,000$ mean casualties for a release of 10 g. Relative to no testing, the sequential testing strategy cuts the number poisoned approximately in half, resulting in tens of thousands of cases. The ELISA testing strategy used in isolation prevents nearly all cases, e.g., if 1 kg is released then the mean number poisoned is 2.82, and six people are poisoned even if the terrorist chooses the worst-case release time within the silo cleaning cycle.

Table 2 contains the results of a sensitivity analysis of isolated changes in 10 key parameters in the no-testing case. Five of these 10 changes impact the number of casualties in the no-detection case (Table 3). Graphs corresponding to Tables 2 and 3 appear in *Supporting Appendix*. The first 4 of these 10 changes involve milk storage and processing. Reducing the time between silo cleanings from 72 to 48 h lowers the number poisoned by $\approx 30\%$ in a large attack with no detection but otherwise has a modest impact. Increasing the silo size from 50,000 to 100,000 gallons (several raw milk silos in California hold up to 200,000 gallons), while varying the number of silos so that the total silo capacity is fixed at 400,000 gallons, and maintaining a dedicated processing line for each silo leads to slightly fewer casualties for small releases but up to twice as many poisoned for large releases and no detection. Similarly, allowing milk from four silos to mix during downstream processing can quadruple the number of casualties in a large attack with no detection. Because the toxin inactivation rate may be very sensitive to the pasteurization temperature and time in the neighborhood of the current pasteurization formula (7), we consider a pasteurization process that causes a 2-log reduction in active toxin. This leads to a huge reduction in casualties if the release size is 10 g or less but has no impact for a 1-kg release.

The remaining six changes are from the downstream portion of the supply chain. We could not find reliable data on the speed of the distribution channel. More rapid distribution leads to

Table 1. Base-case values for model parameters

Parameter description	Value
Production rate	10 gallons per cow per day
Silo size	50,000 gallons
Silos per processing line	1
Time between silo cleanings	72 hr
Speed of distribution channel	80% of milk purchased within 48 hr
Consumers per gallon of milk	4
Time to consume 1 gallon of milk	84 hr
Dose-response probit slope	4.34
Adult ID ₅₀	1 μg
Child ID ₅₀	0.43 μg
Median incubation (adult and child)	48 hr
Dispersal factor of incubation	1.5
Number of symptomatics until detection	100
Time to detect attack is milkborne	24 hr
Testing-detection limit (mouse, ELISA)	60.6 ng/gallon, 303 ng/gallon
Testing-time delay (mouse, ELISA)	48 hr, 3 hr
Fraction of toxin not inactivated by pasteurization	0.316
Fraction of milk consumers who are children	0.25
Fraction of milk consumed by children	0.4

earlier consumption and faster diagnosis, and the former effect appears to dominate, leading to larger attack sizes. Our base-case value for the time to drink a gallon of milk is based on the conservative assumption that everyone has the same consumption rate. However, there is considerable heterogeneity in consumption rates across the population, which causes heavier consumers to buy milk more frequently. Hence, we assume it takes 24 rather than 84 h for a gallon to be consumed. As in the case of rapid distribution, a higher consumption rate leads to more casualties. The dose-response data in Tables 2 and 3 are based on monkey data, which are more plentiful than human data. As in the pasteurization case, the monkey data lead to a drastic reduction in casualties for a small release but have little effect in a large release. Because children rarely eat in restaurants or eat home-canned food, nearly all of the historical incubation data are based on adults. We assume that the median

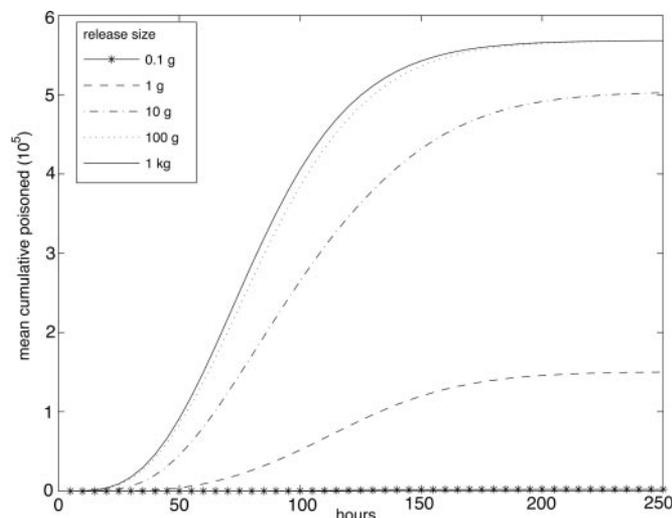


Fig. 2. The mean cumulative number of people poisoned over time for various release sizes in the absence of any detection.

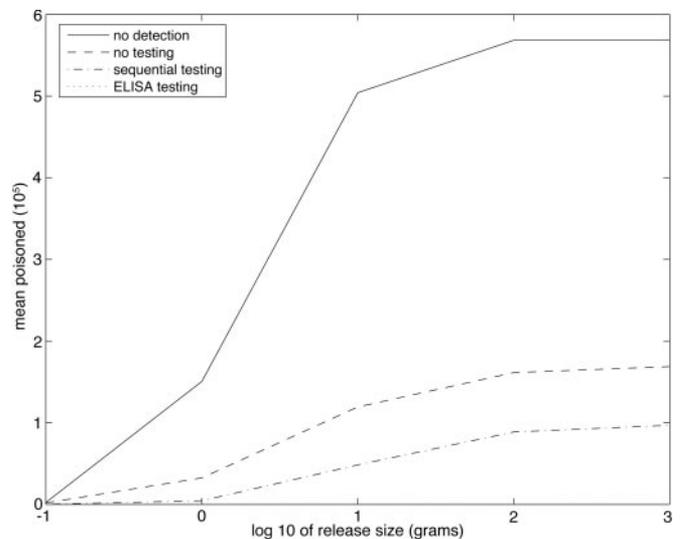


Fig. 3. The mean total number of people poisoned vs. release size for various detection scenarios.

incubation time for children is reduced from 48 to 12 h because of their smaller mass and larger consumption of tainted milk, which lead to earlier detection and many fewer casualties. Our last two changes relate to detection time. The Centers for Disease Control and Prevention maintains a well established national surveillance system for botulism (14) that has been enhanced in the last several years. Botulism in virtually all jurisdictions is an immediately reportable disease, and the characteristic clinical features of botulism suggest that the outbreak might be recognized promptly (e.g., by the presentation of the 10th case). Moreover, because most metropolitan areas have only one or two children's hospitals, and because milk is one of the few staples in children's diets, the time to detect the outbreak as milkborne might be rather quick (e.g., 12 h). Not surprisingly, both changes lead to a reduction in the number of people poisoned.

Discussion

Combating bioterrorism requires an appropriate mix of prevention, mitigation, detection, and response. Our observation that, due to the successive mixing operations in the upstream portion of the supply chain, the impact of a deliberate release upstream of the processing plant is independent of the precise location may aid in prioritizing resources for prevention. A foodborne attack is much more preventable than an airborne or mailborne attack, due to the restricted number of release locations. Requiring all tanks, trucks, and silos to be locked when not being drained or filled would be an obvious step forward, as would security checks for personnel who have access to prebottled milk (farm laborers, truck drivers, receiving labor at the processing facility, and plant engineers) and requiring one person from each stage of the supply chain to be present while milk is transferred from one stage to the next (15). Although these and other measures are included in proposed Food and Drug Administration guidelines (16), they are currently voluntary. Homeland security officials need to engage industry leaders to establish the most appropriate way to guarantee these guidelines are enforced. Although enforcement options range from voluntary guidelines to new laws, the most promising approach may be to develop International Organization for Standardization (ISO) security standards that are analogous to the ISO 9000 standards for quality management and the ISO 14000 standards for environmental

Table 2. Sensitivity analysis for 10 parameters in the no-testing case

Case description	Release size				
	0.1 g	1 g	10 g	100 g	1 kg
Base case	1.7×10^3	3.2×10^4	1.2×10^5	1.6×10^5	1.7×10^5
Time between silo cleanings = 48 hr	2.0×10^3	3.5×10^4	1.2×10^5	1.5×10^5	1.5×10^5
Silo size = 100,000 gallons	2.1×10^2	3.2×10^4	1.6×10^5	2.7×10^5	3.0×10^5
Silos per processing line = 4	4.4×10^1	2.7×10^4	1.9×10^5	4.5×10^5	5.3×10^5
Inactivation by pasteurization = 0.99	6.6×10^{-11}	5.0	1.3×10^4	7.3×10^4	1.5×10^5
Distribution: 90% purchased \leq 24 hr	2.0×10^3	4.5×10^4	1.8×10^5	2.4×10^5	2.6×10^5
Time to consume a gallon = 24 hr	1.9×10^3	6.2×10^4	1.5×10^5	1.7×10^5	1.7×10^5
ID ₅₀ (adult, child) = 70 μ g, 30 μ g	1.8×10^{-16}	6.7×10^{-3}	4.3×10^3	4.2×10^4	1.3×10^5
Median child incubation = 12 hr	6.4×10^2	7.5×10^3	3.4×10^4	5.6×10^4	6.0×10^4
Symptomatics until detection = 10	1.1×10^3	2.0×10^4	8.4×10^4	1.2×10^5	1.2×10^5
Milkborne detection time = 12 hr	1.5×10^3	2.0×10^4	6.9×10^4	9.3×10^4	9.6×10^4

Each change from the base-case value in Table 1 was made in isolation and shown are the mean number of poisoned people computed for five different release sizes.

management (www.iso.ch/iso/en/iso9000-14000/index.html, accessed on November 12, 2004).

Turning to mitigation, botulinum toxin cannot be completely inactivated by radiation (17) or any heat treatment that does not adversely affect the milk's taste. Ultrahigh-temperature (UHT) pasteurization (performed to provide extended shelf life) appears capable of completely inactivating botulinum toxin in milk, but UHT milk has not been embraced by U.S. consumers. Nonetheless, it is worthwhile to perform pasteurization studies to determine whether a more potent inactivation process can be used without compromising nutrition or taste, particularly because the inactivation rate appears to be quite sensitive to the pasteurization temperature and time in the neighborhood of the current pasteurization formula (7). Reducing the time between silo cleanings decreases the number of people poisoned in, at most, a linear manner, but more frequent cleanings would not only increase variable material and labor costs but would possibly require fixed investments in additional silos.

Before discussing detection, we note that, on the response side, $\approx 60\%$ of poisoned individuals would require mechanical ventilation (6). Given the small number of ventilators and limited amount of antitoxin in the national stockpile, the death rate from a large attack would likely be closer to the pre-1950s 60% rate (18) or the 25% rate incurred in the 1950s than to the 6% death rate experienced in the 1990s (19). Moreover, the current treatment, a passive immunization with equine antitoxin, does not reverse existent paralysis, and postexposure prophylaxis with antitoxin has adverse side effects (19). Although an economic impact assessment of this scenario is beyond the scope of our study, the economic cost (including direct medical costs and lost productivity due to illness and death) from a hypothetical botulism outbreak that poisons 50,000 people was estimated to be 8.6 billion (20), using a direct medical cost (assuming ample

ventilators and antitoxin) per hospitalized patient of $\approx \$55,000$ (based on Canadian dollars in 1993–1994). In contrast, two recent U.S. victims receiving injections of “fake Botox” each incurred a \$350,000 medical bill in the first 2 weeks of illness [S. Z. Grossman (lawyer of Botox victims), personal communication]. If this latter amount was spent on each survivor in an attack that poisoned several hundred thousand people, then the total medical costs would be tens of billions of dollars.

Our study highlights the value of rapid in-process testing for detecting an attack, and because stockpiling sufficient ventilators and antitoxin in the event of a large-scale attack would be exorbitantly expensive, it seems wise to aggressively invest in rapid, sensitive, and specific in-process testing. A variety of different botulinum testing technologies are being investigated as alternatives to the mouse assay [summary of the National Institute of Allergy and Infectious Diseases (NIAID) expert panel on botulinum diagnostics, May 23, 2003, www2.niaid.nih.gov/NR/rdonlyres/BB1DDC43-1906-4450-8983-DB0BE3744746/0/bottoxinsmtg.pdf, accessed on November 15, 2004], although published data exist only for the ELISA assay. The current ELISA test appears to be ≈ 2 orders of magnitude more sensitive than needed: if milk in a truck contains 300 ng per gallon, which is the detection limit of the assay (12), the milk gets diluted by a factor of ≈ 20 during processing, and hence each person consumes ≈ 4 ng in their quart of milk, which is 2 logs less than the estimated ID₅₀ for children, using the human data. Therefore, the current test can afford to lose some of this sensitivity if it leads to increased specificity or speed. An alternative less-sensitive ELISA assay based on the catalytic activity of the toxin is also available for botulinum toxin A (21) [List Biological Laboratories (Campbell, CA), www.listlabs.com, accessed on July 1, 2004] and may be more specific in foods (unlike milk) where the toxin is unstable.

Current antibiotic residue testing takes 45 min, during which time the truck waits before having its contents drained into a silo.

Table 3. Sensitivity analysis for five parameters in the no-detection case

Case description	Release size				
	0.1 g	1 g	10 g	100 g	1 kg
Base case	2.3×10^3	1.5×10^5	5.0×10^5	5.7×10^5	5.7×10^5
Time between silo cleanings = 48 hr	2.8×10^3	1.6×10^5	3.8×10^5	3.9×10^5	3.9×10^5
Silo size = 100,000 gallons	2.1×10^2	1.4×10^5	8.4×10^5	1.1×10^6	1.1×10^6
Silos per processing line = 4	4.4×10^1	1.1×10^5	1.2×10^6	2.2×10^6	2.2×10^6
Inactivation by pasteurization = 0.99	6.6×10^{-11}	5.0	3.8×10^4	3.6×10^5	5.7×10^5
ID ₅₀ (adult, child) = 70 μ g, 30 μ g	1.8×10^{-16}	6.7×10^{-3}	7.5×10^3	2.1×10^5	5.3×10^5

Each change from the base-case value in Table 1 was made in isolation, and the mean number of poisoned people was computed for four different release sizes.

A test that takes >45 min is impractical because it either would increase the waiting time for each truck (if milk is not released to the silo until the test results are received) or would need to have near-perfect specificity (if milk is released before the test results are received). In contrast, three possible approaches can be used to deal with a positive result from a sub-45-min test: the truck can be held until a confirmatory mouse assay is performed, the milk can be discarded, or the milk can be routed to a processing line for ultra-high-temperature pasteurization, which kills all of the botulinum toxin. The likelihood that positively tested milk contains toxin may be extremely small, e.g., by Bayes' rule, if there is a 10% probability of an attack occurring in the U.S. over the next 5 years, and the false-positive rate is 10^{-4} , then the probability that positively tested milk contains toxin is only 5×10^{-5} . Regardless of which of the three options is used, it seems clear that a sub-45-min test is necessary from a practical perspective. Even if such a test is not perfectly specific, it could still be an immensely useful tool that could essentially eliminate the threat of this scenario. Even if the total cost of a test was \$50, testing each 5,500-gallon truck would increase the cost of milk by only 1 cent per gallon. In addition, because thousands of people would be poisoned per hour in this scenario, it is imperative to perfect the design and implementation of a near-instantaneous product recall and disposal strategy.

To understand the impact of changing these processing parameters and to assess the danger of bioterror threats to various food industries, we need some understanding of the terrorists' capabilities. To put the release sizes in Figs. 2 and 3 into perspective, we note that the maximum concentration of botulinum toxin in culture is $2\text{--}3 \times 10^6$ mouse units per ml (22), where a mouse unit is the mouse intraperitoneal LD_{50} in micrograms. In the 1980s, the Iraqi bioweapons program apparently increased this concentration 5- to 10-fold with the use of sulfuric acid (23). If so, it would appear that terrorists should be capable of a concentration of at least 3×10^7 mouse units per ml = 4 g per gallon. That is, a terrorist with this technology could easily deliver 10 g of toxin without any special gear. Referring to Fig. 2, in the absence of detection, this amount would poison >400,000 people. Delivering 100 g or more with this technology would be more cumbersome and would greatly increase the likelihood of intercepting the attack. Amplification technologies have advanced significantly in recent years (24), and hence terrorists may be capable of concentrations considerably higher than 4 g per gallon.

Section 5 of *Supporting Appendix* analyzes three additional interrelated issues: secondary cases due to crosscontaminated milk, product tracing, and product recall. Two locations in the supply chain, trucks that are cleaned daily but that make two trips daily and processing lines that are cleaned daily, offer the opportunity for uncontaminated milk to become tainted by uncleaned residue from the primary release. The secondary effect from a release in a truck has an $\approx 50\%$ chance of causing damage equivalent to a release that is 8 h later and $\approx 0.5\%$ as large as the primary release. According to Figs. 2 and 3, secondary casualties would be significant only in cases when the primary release poisons nearly all of its consumers (in the absence of detection). The secondary impact due to tainted processing lines is likely to be much smaller, but the resulting milk concentrations are more difficult to estimate.

This potential for crosscontamination, coupled with consumer anxiety, would probably cause the supply chain's entire milk supply to be recalled and discarded at the time of detection. For the values in Tables 4 and 5, which are published as supporting information on the PNAS web site, this amounts to 4.83 million gallons, which includes 2.24-million-gallon containers of partially consumed milk that need to be recalled from consumers (Eq. 29 in *Supporting Appendix*). In addition, 640,000 gallons per day of freshly produced milk would need to be discarded until the

attack is effectively investigated, the supply chain is turned back on, and consumer confidence returns. This delay could be hastened by effective product tracing, decontamination, and risk communication. The U.S. dairy industry traces every milk carton back to its processing facility, which, at least in theory, prevents the entire nation's milk supply (≈ 300 million gallons) from being discarded and recalled. In other food scenarios where there is no risk of crosscontamination (e.g., fresh produce packaged in the field), the ability to trace a product back through the particular path it takes in Fig. 1 could lead to a significant reduction in the amount of product recalled and discarded. As an illustration, we compute (Eq. 30 in *Supporting Appendix* and Table 6, which is published as supporting information on the PNAS web site) the amount of milk that needs to be discarded as a function of the release location (farm, truck, or silo) and the stage (cow, farm, truck, silo, or processing facility) to which the milk can be traced, hypothetically assuming no crosscontamination.

Our sensitivity analysis suggests there are three types of variables. Variables of the first type (time between silo cleanings, silo size, and number of silos per processing line) cause a vertical shift in the number poisoned vs. release size graphs (Fig. 5 *a-c*), which is published as supporting information on the PNAS web site) and underscore the subtle relationship between high production efficiency and the consequences of a bioterror attack. Economies of scale can represent a double-edged sword: increasing the time between silo cleanings, silo size, or number of silos per processing line increases the amount of contaminated milk but reduces the toxin concentration of this milk, thereby mitigating the impact of a small release and exacerbating the effect of a large release. However, for the parameter regimes considered here, the reduction in casualties in a small release is very modest, whereas the increase in casualties in a large release with no testing and poor detection is in the hundreds of thousands. Variables of the second type (ID_{50} , pasteurization inactivation) result in a horizontal shift in the number poisoned vs. release size graphs (Fig. 5 *d* and *g*). More precisely, to cause equivalent damage, the release size for the monkey ID_{50} s needs to be 70 times larger than the release size for the human ID_{50} s. Similarly, to generate an equivalent casualty level, the release size in the 99% inactivation scenario needs to be $1-0.684/1-0.99 = 31.6$ times larger than the release size in the 68.4% inactivation scenario. Variables of the third type (distribution speed, consumption rate, children's incubation, number of symptomatics until detection, and milkborne detection time) all relate to the speed of various events and have no impact on the casualty level if the attack is not detected. In the no-testing case, the resulting graphs (Fig. 5 *e-f* and *h-j*) are very similar to one another and, for the parameter values considered here, the change in the children's incubation has the biggest impact, and the consumption rate has the smallest impact.

Conclusion

In closing, it is important to stress that several elements of the model contain enough irreducible uncertainty to preclude estimating the impact of an attack to within several orders of magnitude. First and foremost is the dose-response curve. The paucity of human data makes an estimate of the ID_{50} a difficult task, and a reliable estimate of the probit slope is impossible. The ID_{50} values used here are not close to the worst-case estimate, due to the possibility that several sublethal (injected or oral) doses collectively containing 1-10% of the LD_{50} may be lethal, as in guinea pigs, rabbits, and mice (25). There are also three aspects of the model that have not been discussed in the open literature, although presumably studies can and perhaps have been performed: the inactivation rate attained by pasteurization, the specificity of an ELISA test in milk, and the release size that a terrorist organization is capable of. Such studies would allow our results to be sharpened considerably. The dose-response

