Defining a Lowest-Risk UAS Category

Walter Stockwell, Ph.D.
Brendan Schulman, J.D.

DJI Research LLC

DJI Research, LLC

435 Portage Avenue

Palo Alto, CA 94306

walter.stockwell@dji.com

brendan.schulman@dji.com

December 9, 2016

Abstract

Not every unmanned aircraft system (UAS) poses the same level of risk. Intuitively, the smaller and lighter the UAS, the less risk it poses. Regulators are increasingly finding the need to declare a certain category of small unmanned aircraft systems (UAS) to be “lowest-risk” or relatively “harmless.” This categorization allows them to set aside the devices that require the simplest regulation while maintaining an acceptable level of safety, and focus rulemaking and technical effort on the higher-risk UAS. An early and hurried effort to define such a category in regulation, by the United States Federal Aviation Administration Registration Task Force (“RTF”), has seemingly unintentionally influenced other jurisdictions to follow suit in a broader fashion without close analysis of the underlying analysis or assumptions. We show that the method of estimating the safety of a UAS used in initial estimates proposed by the RTF reveals that safety is dominated almost entirely by factors unrelated to the physical characteristics of the UAS, such as pilot skill and exposed population. We further show that current academic and regulatory predictions of injury based on UAS kinetic energy and relied upon by the RTF are unrealistically conservative. The combination of these factors reveals that the RTF was overly conservative in its selection of a 250 gram weight threshold, given the framework it chose to define the category, and that this weight limit was likely too low by an order of magnitude. A more realistic upper threshold for a lowest-risk mass-based category of UAS is found around the 2 kilogram range.

Keywords: UAS, safety, harmless, risk, policy

Defining a Lowest-Risk UAS Category

Unmanned aircraft systems have become ubiquitous around the world. The United States since late 2015 has registered more than 600,000 small UAS owners, with the FAA estimating that each registrant owns three UAS. Accounting for people who have not yet registered or who are exempt, there are likely 3 million or more civil UAS in operation in the United States. This is larger than the entire manned aircraft fleet in the United States. Similar numbers are expected in the European Union[1] and Asia. Most of these vehicles are relatively small (<=2.5 kg) and most are multirotors with four or more propellers. The huge surge in interest in these aircraft and the sudden realization that there are millions of them flying right now has caused regulators worldwide to try to quickly formulate workable regulations for these systems. These regulations must balance safety, accessibility, and regulatory burdens. Many regulatory agencies have desired to define a category of small, "lowest-risk" UAS that would be subject to little or no regulation.

The FAA Registration Task Force

At the end of 2015, the FAA convened a task force to quickly recommend an approach to a new registration system for small UAS in the United States. The Unmanned Aircraft Systems (UAS) Registration Task Force ("RTF"), essentially an aviation rulemaking committee composed of 25 members from a cross-section of aviation stakeholder groups,[2] was given less than thirty days to deliberate and deliver a recommendation report to the FAA. One key RTF recommendation involved establishing a category of UAS that were considered low-risk and

---

[1] The EASA/NAA Task Force Report Study and Recommendations regarding Unmanned Aircraft System Geo-Limitations estimated there are 3 million small UAS in the EU based on industry surveys.

[2] DJI served as a member of the RTF.

would not require registration. The RTF chose a cutoff based on mass, and recommended that UAS weighing less than 250 grams should be exempt from registration requirements. As stated in the report, "[t]his approach best satisfied the Task Force's concerns about safety and provided a minimum weight threshold for registration that is easy to understand and apply and would therefore encourage compliance."[3]

The safety concern was addressed by finding a mass limit that met an overall policy goal of keeping expected risk of UAS as safe as the experienced risk of manned aviation. The RTF used a three-part approach to estimating risk: a standard risk equation taken from very old estimates of risk at military missile ranges; an estimate of probability of fatality (PoF) vs. kinetic energy (KE) of an impact taken from very old estimates of lethality of shrapnel in explosions; and a standard calculation of the terminal velocity of a generic falling object.

The RTF started with an old rule-of-thumb that an impact involving 80 Joules of kinetic energy marked a threshold between less dangerous and more dangerous impacts. As the RTF noted in its report, "an object with a kinetic energy level of 80 Joules (or approximately 59 foot-pounds) has a 30% probability of being lethal when striking a person in the head."[4] In other words, the RTF found it socially acceptable and acceptable as an aviation policy matter to leave essentially unregulated (or lightly regulated) the UAS devices that have long been assumed to pose a 30% lethality risk in the event of an impact with an unprotected person.

However, the RTF desired to determine a weight representing the same risk, as weight is a convenient method to categorize a UAS. The key to finding the weight threshold lay in a probability

---

[3] Unmanned Aircraft Systems Registration Task Force Aviation Rulemaking Committee (UASRTFARC) Final Report, page 6

[4] UASRTFARC Final Report, page 8

model that equated a human impact of 80 J of kinetic energy with the 30% probability of fatality. This KE/PoF curve was sourced to MITRE, and looked like this:

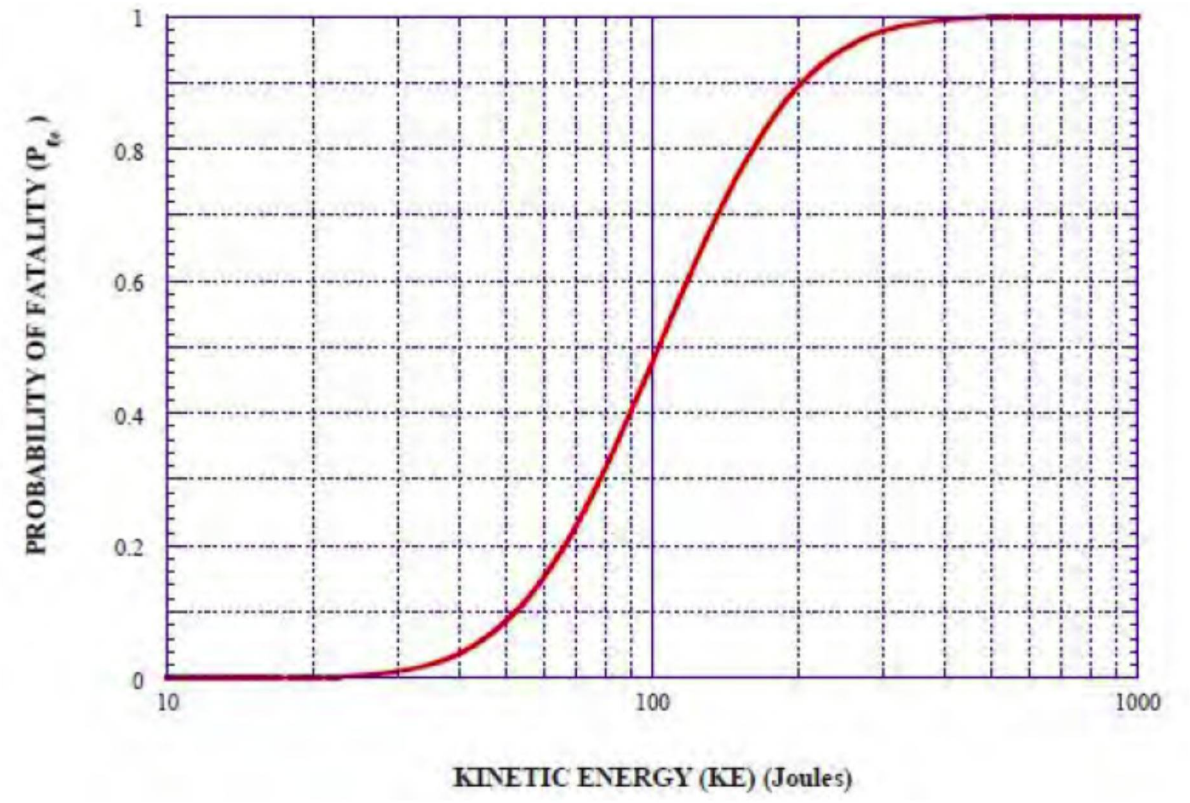

Figure 1. Kinetic Energy vs. Probability of Fatality[5]

The result of the use of this model, plus other estimates, assumptions, and calculations aimed at estimating the kinetic energy of a falling UAS, resulted in the 250g weight limit.

Notwithstanding the inclusion of calculations in its report, the RTF was careful to expressly document that this weight calculation was not a rigorous determination, was reached as part of a stakeholder consensus process and was “interdependent on the Task Force’s other recommendations

---

[5] This figure was referenced by the RTF as the source of their PoF values. It appears in Lacher and Maroney, 2012 where it is referenced to Henderson, 2010.

on the registration process."[6] The report noted that some members "felt there was insufficient time afforded to fully evaluate the calculations and assumptions made that resulted in the 250 gram cutoff weight."[7] Given the circumstances, it is not surprising that all members of the Task Force *unanimously* emphasized the limited nature of its recommendations and warned against using the conclusion for any other purpose: "It should also be noted that the 250 gram weight threshold was agreed to for registration purposes only and was not a validation of the underlying assumptions for any purpose other than the registration requirement. It was agreed by all members that this threshold, arrived at under the circumstances described, should not be used by the FAA to establish operational restrictions or categories in any future rulemaking."[8] The recommendations in the report were offered to the FAA as a "holistic package, with elements of each recommendation closely interconnected with the others." The RTF combined the low mass limit of 250g with other recommendations intended to encourage and facilitate participation in the registration process such as the ability for one person to register multiple aircraft under the same number, and the recommendation that registration be online and free of charge.[9] Subsequently, and as intended, the FAA promulgated new registration regulations that adopted the recommended 250 gram (0.55 pound) threshold for registration, and virtually all of the other recommendations as well. *See* 14 CFR 48.15(b).[10]

---

[6] UASRTFARC Final Report, page 9

[7] UASRTFARC Final Report, page 9

[8] UASRTFARC Final Report, page 10

[9] The intent of this paper is not to question the consensus-based package of recommendations made by the RTF within the context of its designated task concerning a simplified registration system, but rather to examine whether the 250-gram mass identified by that committee is a meaningful threshold for other regulatory determinations.

[10] 14 CFR 48 Registration and Marking Requirements for Small Unmanned Aircraft, www.ecfr.gov

Notwithstanding the clear and explicit qualifications stated by the RTF, regulators outside the United States have misunderstood the RTF recommendation and given it more significance than it deserves, especially with regard to the proposition that 250g represents a significant threshold between a "lowest-risk" UAS and a more dangerous UAS requiring comprehensive regulations. The 250g cutoff is seemingly the first time that particular mass has been used for any aircraft regulation anywhere in the world. Subsequently, however, a 250g threshold has been referenced in proposals around the world, including in EASA's recent Prototype Regulation for drones,[11] and in news reports from Canada concerning proposed Transport Canada regulations expected next year,[12] and in Russia.[13] In a similar manner, Austria has used 79 J KE as way to distinguish between UAS that need regulation and those that can be safely ignored.[14] The 250g weight threshold recommendation is thus having an influence upon UAS regulations in ways that go far beyond the original intent, raising the prospect of regulatory approaches that do not appropriately balance actual risk with the benefits of use, cost and regulatory burden.

We will examine the safety estimate with more scientific rigor in this paper, in a way that might have been conducted by the RTF if it had not been given such a limited amount of time to reach a negotiated, non-scientific, consensus-based recommendation for the FAA.

---

[11] *See* European Aviation Safety Agency, *Prototype' Commission Regulation on Unmanned Aircraft Operations*, Aug. 22, 2016, *available at* https://www.easa.europa.eu/system/files/dfu/UAS%20Prototype%20Regulation%20final.pdf .

[12] *See* "Rigorous rules proposed for recreational drone flyers, documents show," CBC News, October 19, 2016, available at http://www.cbc.ca/news/canada/ottawa/transport-canada-drone-regulations-1.3810123 .

[13] *See* "Russia's New Drone Rules Look a Lot Like America's," Popular Science Jan. 4, 2016, available at http://www.popsci.com/russias-new-drone-rules-look-lot-like-americas .

[14] The Austrian Aviation Act, § 24d. Unbemannte Geräte bis zu 79 Joule maximaler Bewegungsenergie

The Risk Equation

Most analyses of "lowest-harm" approach the subject with a risk calculation. Both the FAA and EASA have stated as a policy goal that UAS should be no more dangerous than manned aviation. This implicitly sets an expectation of an allowed number of casualties per flight hour. For example, in general aviation, the current accepted safety goal for aircraft design is a casualty rate of 10^{-5} fatal accidents per flight hour.[15] According to the Airplane Owners and Pilots Association (AOPA), the current general aviation accident rate is approximately 1.33 fatal accidents per 100,000 flight hours.[16] Regulators and the aviation community always strive for safety improvements, but this rate appears to be socially acceptable.

The following equation is an example used in the United States for estimating risk from UAV flights in national test ranges, given in the guidance document *Range Safety Criteria for Unmanned Air Vehicles, Rationale and Methodology Supplement*, supplement to RCC 323-99[17]:

$$CE = PF * PD * AL * PK * S$$

Where CE is the expected casualty rate; PF is the UAS probability of failure; PD is the population density at risk from the UAS flight; AL is the lethal area of the impact; PK is the probability of casualty (fatality); S is a sheltering factor. This is the same equation used by the RTF, but with different labels on the factors.

---

[15] See for example AC23-1309.1E, figure 2. Failures considered hazardous, that could result in loss of life, should have an occurrence rate of 10^{-5} for small single engine aircraft.

[16] See General Aviation Safety Record - Current and Historic, available at https://www.aopa.org/about/general-aviation-statistics/general-aviation-safety-record-current-and-historic .

[17] Range Safety Criteria for Unmanned Air Vehicles, Rationale and Methodology Supplement, Supp. RCC 323-99, Appendix D published by the Range Commanders Council.

We can examine this equation by looking at each factor; what are likely ranges for each factor; and how realistic each factor might be. RCC 323-99 calls this risk estimate a method for making a first try at a conservative estimate; if the result is satisfactory to the policy goal, then no further analysis is needed. If the result is not satisfactory, then a more detailed analysis should be performed.

1. Casualty Estimate, CE

The result of the calculation is a casualty estimate in fatalities per flight hour. This is the policy goal and is usually set by the operating history of manned aviation and social acceptance in light of benefits and other factors. There are a few choices here: some calculations will look at total casualties (fatalities plus serious injuries), some will look at fatalities only. Manned aviation will have a relatively high ratio of fatalities to injuries because the crash of an airplane is usually fairly energetic, with lives onboard directly at risk, and therefore more likely to be fatal. Small UAS crashes may have a very different ratio of fatality to injury because there are no people aboard the UAS by definition. As of late 2016 there have been no reported fatalities resulting from impacts with the kinds of small UAS in popular use today, notwithstanding tens of millions of estimated flight hours by such small UAS around the world.

One important choice in the analysis is whether to consider all deaths and injuries caused by aviation or only consider uninvolved people on the ground. RCC 323-99 states the goal of $CE = 10^{-6}$ fatalities per flight hour should be considered as risk on an individual basis.[18] This document refers to an analysis of risk from a different RCC document, RCC 321-00, that looks at

---

[18] Range Safety Criteria for Unmanned Air Vehicles, RCC 323-99, §2.2.1 Casualty Expectation. “Must be less than one casualty in a million flight hours.”

risk from rocket and missile launches[19]; RCC 321-00 examines risks from other comparable activities and industries to justify this number as an acceptable level of risk to the public. Other policy analyses have considered numbers ranging from 10^{-5} (comparable to fatalities in general aviation) to 10^{-8} (strictly limited to fatalities to uninvolved persons from aviation as a whole). We see here a range of three orders of magnitude as an acceptable safety goal for various aviation activities that present a fatal risk to people in the air and on the ground.

Probability of Failure, PF

Probability of failure is the chance per flight hour that a given flight will end in a crash. This includes accidents caused by pilot error and by vehicle failures. The RTF assumed a failure rate of 10^{-2} hr^{-1}as a guess.[20]

A modern electric multirotor, the most common UAS flying in the national airspace, is essentially a consumer electronics product, and may be expected to have a similar failure rate as common consumer electronics. If we look at other types of electromechanical products with an expectation of consumer-grade quality we find failure rates between 5 10^{-4} $hr^{-1}$for printers and 10^{-5} $hr^{-1}$for disk drives.[21] We should expect that even consumer-quality UAS have a similar failure rate, which is two to three orders of magnitude more reliable than the RTF assumption.

The overall UAS failure rate will include pilot errors however, and this rate is harder to quantify. If the product failure rate is much lower than the pilot error rate, then pilot errors will dominate this number. 10^{-2} hr^{-1} might be a plausible guess in this case, with estimates ranging from perhaps 10^{-1} hr^{-1} to 10^{-3} hr^{-1}. This is an interesting result because this implies the quality of

---

[19] Common Risk Criteria Standards for National Test Ranges, RCC 321-00, §3 Risk Criteria Rationale

[20] UASRTFARC Final Report, page 9

[21] https://src.alionscience.com/pdf/TypicalEquipmentMTBFValues.pdf

the aircraft does not drive the safety outcomes – pilot skill and expertise are drivers instead. This is not so different from manned aviation. The most common causes of accidents in general aviation are loss of control and controlled flight into terrain.[22]

Also, it is much harder to identify a UAS event that might be considered a "failure." A bumpy landing of a quadcopter that causes it to flip it over upon landing seems like a non-event (with zero chance of injury). One would not say the same for any manned aircraft. Most UAS incidents are likely not even analogous to fender-benders, but something more akin to dinging a car door in a parking lot. For purposes of this analysis, we will maintain the PF assumption chosen by the RTF, but note that the RTF assumption is likely too high.

Population Density, PD

Population density is expressed as people per square meter and for simplicity is typically divided into categories that roughly correspond to real-life scenarios such as "sparsely populated", "suburban," or "dense urban." Values may range from 10^{-2} m^{-2} (e.g. New York City) to 10^{-6} m^{-2} (e.g. Wyoming, a very sparsely populated state). Population density may be deceiving however. As an example, in New York City, a very dense urban area, population density is closely linked to S, the sheltering factor. Even in a city like New York City, or many European cities with more pedestrians than typical American cities, a high percentage of the population is sheltered at any given time. In the United States, suburbs do not have many pedestrians and the sheltering factor is also relatively high. Almost all the population at any given time will be sheltered inside buildings or vehicles. Unlike a manned aircraft, a typical

---

[22] The General Aviation Joint Steering Committee pareto, http://www.gajsc.org/gajsc-pareto/

small UAS is not heavy enough to penetrate building roofs, automobiles and other shelters so as to expose people inside to risk.

In addition, most UAS regulations require operators to fly within VLOS, away from and not over, unprotected persons. *See, e.g.*, 14 CFR 107.39. Small UAS have a relatively short flight time limited by battery capacity, and are generally used for operations close-in to the operator. So, even in a densely populated *area*, the sUAS may be flown over the localized area of a park with an effectively very low population density. A more realistic estimate, if needed, would require looking at the actual operation to understand what is the actual population exposed to risk from the UAS operation. The RTF assumed a population density of $4\ 10^{-3}$ m-^{2} and a sheltering factor of 0.2. This might be thought of as noon in a typical American suburb, where some number of people will be walking to and from lunch.

Shelter Factor, S

The shelter factor describes the percent of the population that is actually exposed to risk from the sUAS at any particular time. The shelter factor should be considered in conjunction with the population density and will be directly related to the type of environment being overflown. A beach will have a shelter factor close to 1, meaning all of the population is exposed. A dense city however, will likely have the vast majority of the population sheltered at any given time, and might have a shelter factor <0.1. The shelter factor will vary with time as well. At noon, many might leave work or home to shop or eat; at midnight almost everyone will be inside. Any attempt to create population density scenarios as described above ("sparsely populated" vs. "dense urban") should include an estimate of the percentage of the population that is actually sheltered as part of the estimate. We will consider the S in the range of 1 to 10^{-1}, but not tied specifically to population density as a more realistic estimate should.

Lethal Area, AL

Lethal area can be thought of as the cross-sectional area of the UAS as it crashes. Some will consider the average size of a person and increase AL by considering *any* intersection of a person and the UAS as in the lethal area. For high velocity impacts, some analyses will even allow for skipping and multiply AL by some factor to account for multiple impacts. This is common when the analysis is drawing from experience with explosions, large airplane crashes or failed missile launches with resulting shrapnel.[23]

These assumptions are probably not applicable to small UAS, and certainly not to small multirotors. Small UAS will generally not crash at high velocity, unless perhaps they are falling straight down at terminal velocity. Most airframes are plastic, foam or carbon fiber and break upon hard impact, preventing skipping. Without combustible fuel, they will not explode or erupt into flames. A multirotor, the most common UAS today, will not glide, so worst case crashes will either be horizontal at the multirotor maximum speed (typically, pilot error) or vertical at terminal velocity (typically, vehicle malfunction at maximum altitude.) Overly large estimates of AL are therefore not justified. Plausible generous estimates for AL for small UAS may vary from 10^{-1} m^2 (30 cm x 30 cm) to 10 m^2 (3 m x 3 m). The RTF report used the value 2 10^{-2} m^2, which is exactly the cross-sectional area of the small UAS the committee considered to be falling. There was no assumption that the size of the lethal area should be increased due to skipping or a gliding path. This is an appropriate assumption for a sUAS that is assumed to be falling at terminal velocity from some height.

---

[23] See for example Ball, Knott & Burke, *Crash Lethality Model*, NAWCADPAX/TR-2012/196, Dept of the Navy.

Probablity of Fatality, PK

Estimate of the probability of fatality (PoF) is one of the most complex parts of the risk calculation. People often conflate probability of fatality and probability of casualty, which includes serious injury as well as fatality. This question of what is being considered is, first, a policy question, and starts with the question of what level of risk is acceptable. The RTF used the rule-of-thumb of 80 J, which was understood to represent a 30% PoF, as the acceptable goal for a lowest-risk UAS category not requiring registration. Once this is decided one must find a model of impact vs. effect to use in the estimate. This is usually taken to be a function of the kinetic energy at impact of the UAS. Kinetic energy is a function of UAS mass and velocity. A more sophisticated estimate may take into account other factors that may affect injury (such as rigid vs. compressible structure) or details of the ability of the impact to transfer energy (such as a solid ball hitting vs. an extended, flexible structure hitting.) In any case, the highest that this number in the equation can be is 1 (*i.e.* every impact kills); a plausible lower estimate used is often 10^{-2} (one in one hundred impacts will kill). We will examine this in more detail later in the paper. Because this value cannot vary much, and is not more than 1, this factor -- perhaps surprisingly -- does not play a significant part in the risk assessment.

Estimate Results

We can use the range for each parameter as discussed to create a range of plausible estimates. The results are shown in the table below. The values chosen by the RTF are shown for reference.[24]

---

[24] The result shown here is slightly different from the actual RTF value of 4.7 10^{-8} due to rounding. The difference is not significant for this sort of estimate.

Factor	Optimistic	RTF	Pessimistic	Units
PF	10^{-3}	10^{-2}	10^{-1}	hr^{-1}
PD	10^{-6}	$4\ 10^{-3}$	10^{-2}	m^{-2}
AL	10^{-1}	$2\ 10^{-2}$	10	m^{2}
PK	10^{-2}	0.3	1	
S	10^{-1}	0.2	1	
CE Result	10^{-13}	$5\ 10^{-8}$	10^{-2}	hr^{-1}

The optimistic case can be thought of as an experienced pilot flying a small UAS in an empty park; the pessimistic case looks more like a novice pilot flying a larger UAS over a busy summer beach. The results span an extraordinary range of eleven orders of magnitude. The optimistic case is obviously safe; the pessimistic case would not meet the policy goals considered. Perhaps most interesting when considering what would constitute a "lowest risk" UAS however is that PK, the probability of fatality has such a *small* overall effect on the result. Clearly, this *type* of risk calculation cannot be used to decide what is a "lowest risk UAS"

A Closer Look at PK

The probability of fatality or casualty is directly related to the UAS itself. The method used in the FAA Registration ARC Report is typical for estimates drawn from the missile test range and explosion safety communities. The PK estimate breaks into two parts: (1) a way to estimate the kinetic energy of an impact; and (2) a way to relate kinetic energy (KE) of an impact to a probability of fatality (PoF). This same method is used in the various Range Commander Council guidance documentation.

Estimating Kinetic Energy

Kinetic energy estimates for UAS impacts have typically assumed a terminal velocity fall as an assumption of a worst-case scenario. Terminal velocity of a falling object can be calculated using the following equation:

$$V_t = \sqrt{\frac{2mg}{\rho A C_d}},$$

where V_t is the terminal velocity of the falling object, m is the mass of the object, g is the acceleration due to gravity, ρ is the air density, A is the object's cross-section, and C_d is the drag coefficient. Kinetic energy is proportional to V_t^2; therefore one can see that the KE of a falling object is ~m and $\sim(A \times C_d)^{-1}$. Other quantities remaining fixed, an object will have a higher terminal velocity with a larger mass; it will have a lower terminal velocity with a larger cross sectional area or a larger drag coefficient. The only one of these quantities that can be easily measured is the mass. Cross-sectional area is a complicated function of the geometry of the UAS and the actual angle and behavior during falling. The drag coefficient will be a function of the shape and angle of fall, as well as a function of the remaining propellers on the UAS. Given the lack of any time to conduct a study or collect data, the RTF simply assumed a Cd of 0.3. However, this is unrealistically low for a multirotor, the most common form of small UAS in the airspace. In fact, the RTF-selected C_d of 0.3 approximates the drag coefficient of a sports car, a shape that is deliberately streamlined by design to reduce drag. In addition to the bare geometry of the drone fuselage, a falling UAS will have propellers creating extra drag. The most popular model UAS, the quadcopter, has four propellers that are substantially large compared to the size of the entire aircraft, arranged in a cross-arm formation that is not at all aerodynamically streamlined. A sphere has a $C_d \sim 0.5$; a cube has a $C_d \sim 1$. Until actual measurements are made of falling UAS, a value of $C_d = 1$ seems far more reasonable than 0.3.

The RTF report used the terminal velocity equation to calculate the weight of a falling UAS with 80 J (assumed to be 30% fatal), given a $C_d = 0.3$, and that result was approximately 250 g. If one recalculates this with a more realistic $C_d = 1$, the KE of a 250 g UAS drops to about 25 J. *Simply by changing the drag coefficient to something more realistic, the RTF could* *have justified setting a registration cut off at 450g with the same 30% PoF threshold that it had found acceptable.*

Moreover, because, at 450g, we are dealing with a larger UAS than one weighing 250 g, it is reasonable to use a larger estimate of cross sectional area. Increasing the cross-sectional area from 0.02 m^2 to 0.03 m^2 would allow us to increase the UAS mass to 550 g and still meet the 30% PoF threshold that corresponded to the safety goal set by the RTF at the outset. The Parrot Bebop 2, a popular small consumer quadcopter, has a mass of 500g and is roughly 32 cm x 28 cm. The kinetic energy of the 500g Bebop 2 might be expected to be only 45 J at terminal velocity, well below the 80 J limit used by the RTF. Clearly, 250g was not the right mass given these "reality check" adjustments to the RTF's work. The adjustments easily suggest a "correct" number of more than double the mass.

In addition, it is improbable that a UAS is able to efficiently transfer its KE to a target. Because UAS are in the form of extended objects, not compact balls, they will tend to rotate as they impact an object. This rotation prevents the full amount of energy from being transferred in an impact. Preliminary research suggests that only a fraction of the UAS kinetic energy, perhaps as little as one third, will transfer to a person during an impact. In contrast, the Registration Task Force simply assumed that 100% of kinetic energy would transfer from the UAS in the impact. The assumption of 100% energy kinetic energy transfer assumes that a UAS impact will be like a bullet or shrapnel impact. This is clearly unrealistic and inappropriate for small UAS. If we assume that only 33% of the kinetic energy is transferred in a typical impact, we could conclude that a 1.1 kg UAS in free-fall at terminal velocity meets the RTF's safety standard of an 80 J transfer energy.

Kinetic energy alone will not determine a safe weight, as the weight one calculates is very sensitive to the assumptions about cross-sectional area, drag coefficient, and efficiency of transferring energy during impact. In its haste, the RTF had to use shortcuts and assumptions that are not well-matched to the characteristics of a UAS. By making this additional "reality check" adjustment to the calculation, we can conclude that the RTF's calculation of mass for its cutoff should have been over four times higher.

Probability of Fatality

As mentioned, in order to merge its policy-oriented safety goal of acceptable PoF to one that defined a category of UAS by mass, and to begin its work in defining a lowest-risk category, the RTF first made a determination of the kinetic energy that would represent a potentially fatal threat to a person in an impact. To do this, the RTF relied on a graph of a curve of KE vs. PoF found in a 2012 MITRE report on UAS safety.[25] But this graph is not nearly as recent as it seems, and depends on assumptions that are invalid in the modern world.

The PoF graph (shown previously in Figure 1) was taken directly from a UK Ministry of Defense report, written by Deputy Chief Inspector Explosives Jon Henderson, on lethality of debris fragments from accidental explosions.[26] Henderson's paper, cited by the RTF report, describes some of the history of this curve (which he did not create) and other ideas about lethal impacts, such as the legendary notion that an 80 J impact is often fatal. Henderson presents a

---

[25] "Referencing information from a 2012 MITRE report (which further references a United Kingdom Ministry of Defense 2010 study), an object with a kinetic energy level of 80 Joules (or approximately 59 foot-pounds) has a 30% probability of being lethal when striking a person in the head." UASRTFARC Final Report page 8.

[26] Henderson, J., *Lethality Criteria for Debris Generated from Accidental Explosions*, UK Ministry of Defense, 2010.

few different curves estimating KE vs. PoF, and compares them to everyday objects, such as ball sports (e.g., cricket, golf, baseball, tennis.) Henderson is quite clear that the graph at issue should be considered very conservative, as it presents a much higher PoF for a given KE than other curves published in the explosives literature, and especially when compared to everyday experience with sports such as cricket and baseball. Cricket and baseball involve unprotected players exposed to balls with a typical KE > 140 J.

The lethality curve in Henderson traces back one step to a U.S. Department of Defense Explosives Safety Board (DDESB) report from 2007.[27] The DDESB report presents the curve as a composite of lethality data taken from the Range Commanders Council report of 2000, RCC 321-00.[28] The Range Commanders Council prepares regular updates to their reports on calculating risk at the White Sands Missile Range, with separate specific reports for inert debris and UAVs. Both the inert debris and the UAV reports use the same KE vs. PoF curve, though the report on UAVs, RCC 323-99, comments that this curve is taken from RCC 321-00 which was designed for ballistic missile impacts.

The supplement to RCC 323-99, in turn, advises to just assume the PoF = 1 for simplicity, with the possibility that "Exceptions might be for debris from very light weight material UAVs."[29] This is consistent with the idea that the actual PoF for a UAS impact has little in common with a safety analysis dealing with exposure to an explosion or large crash.

---

[27] Swisdak, Tatom, & Honig, *Procedures for the Collection, Analysis and Interpretation of Explosion Produced Debris*, Department of Defense Explosion Safety Board, DDESB TP 21, 2007.

[28] RCC 321-00, 2000. Figure 4-1 presents separate KE vs. PoF curves for impacts to the head, thorax, and abdomen or limbs.

[29] RCC 323-99, page D-4

Critically, we are now at least four layers deep in the references (RTF citing MITRE, citing Henderson, citing DDESB, citing RCC) but none of those reports did any independent research or evaluation of the risks in order to create or validate the safety curve. The curve has been merely passed from one report to the next for almost 20 years. Of course, none of these reports was written for the solemn purpose of creating legal requirements or restrictions on the use of technology within a category.

The Range Commanders Council report 321-00, finally, points to the data behind the curve. That data comes from a 1968 report by a group of researchers who documented a computer code made to estimate the effect a wide-spread *nuclear war* would have on the United States.[30]

The nuclear casualty report itself relied on previous studies involving shooting balls and other impactors at animals, force to crush cadaver skulls, and other actual data relating various kinds of trauma to injury. The authors realized these data had a basic problem, however. Most of the studies for impact, for example, had been made using fairly light, compact objects, like small steel balls or cubes. Their study needed to extrapolate from the data they had to the effects of much larger impacts, such as might happen when a building is blown into a person's body during a nuclear blast. So the authors created curves with data they had, which involved data from impactors less than 10 grams in mass, extrapolating to estimate the effect from impactors ranging from ~1 gram to ~50 kilograms. Their simulation needed to account for very small impacts from shattered windows, up to big impacts from chunks of concrete buildings. Interestingly, they thought their extrapolation was reasonable because it showed rough

---

[30] Feinstein et. al., *Personnel Casualty Study*, ITT Research Institute, 1968.

agreement with the old "59 ft-lb" (80 J) rule of thumb. The authors describe the situation they faced compiling this data: "Unfortunately, very little information was available relating specifically to mortality. Consequently, many judgments were made, rendering the results qualitative… The task for future experimenters will be to gather the types of data required for casualty estimating."[31]

The authors made additional assumptions about mortality in their estimates. Because they were looking at the consequence of a full-scale nuclear war, they assumed the civil infrastructure would be destroyed, specifically including the civil heath care system. Therefore, they assumed that almost any wound could be considered fatal. So, for example, they included in their calculation that "superficial" injuries would carry a 10% fatality rate. "Therefore, the effects of severity and probability of occurrence for each effect have been combined, averaged and extrapolated as necessary to obtain one continuous range of missile masses and velocities which might be of interest where data were available. For example, … the effect of penetrating glass is classified as a superficial wound, these are estimated at 10 percent mortality."[32] Because they were making these estimates for a very specific case (a full nuclear exchange) the authors created their models to match that specific case, with assumptions that are not appropriate for general injury models in modern society. The assumption that every laceration is significantly lethal is blatantly absurd even in the modern world of 1968, let alone a world now advanced enough to enable people to operate their own personal flying robots.

This report, at best, documents a computer program that can be used to make a fairly conservative estimate of the effect of a full-scale nuclear war on an industrial society. Yet this

---

[31] Feinstein, et. al., page 21
[32] Feinstein, et. al, page 27

analysis is the source of the now-legendary KE vs PoF curve underlying all notable research on UAS impact risk for the 48 years that followed. This fact alone makes efforts to define a PoF threshold based on kinetic energy suspect from the start, and makes the any reliance on the curve for expansive UAS policymaking appear ill-conceived.

We can find some real-world data that informs our view of what a realistic curve of PoF might be. An interesting look at battlefield PoF was made during the Korean War. The experience of World War II was fresh in the minds of the US Army, and Army researchers were beginning to take a close look at the effectiveness of the weapons used in WWII. In 1952, the Operations Research Office of Johns Hopkins University presented a classified study to the U.S. Department of Defense examining requirements for a modern battle rifle.[33] This study was declassified in 1972. This report was one of the arguments that ultimately led to the adoption of the M16, a rifle that fired a lighter .223 caliber bullet compared to the older .30 caliber bullets used in WWII. The report first looks at the operational history and lethality of the rifle as a weapon system in WWII and concludes that the .30 rifles used in WWII have a "lethal index (ratio of kills to hits) exceed[ing] 30 percent."[34] The report notes that the data included all injuries from rifle bullets, with hits to random parts of the body. The report also notes that almost all injuries occurred at battlefield ranges less than 300 yds. This is therefore a direct measurement of a PoF vs KE for random hits to the body, exactly what the graph used by the RTF purports to be. A .30 caliber bullet fired from a typical WWII rifle will have more than

---

[33] Hitchman, Norman, *Operational Requirements for an Infantry Hand Weapon*, Operations Research Office, Johns Hopkins University, 1952.

[34] Hitchman, page 8.

4000 J at the rifle muzzle, and more than 2000 J at 300 yds.[35] Therefore, actual battlefield data would calculate a 30% PoF impact only at an energy value *at or above 2000 J*, not at the 80 J level selected by the RTF. (If 2000 J were selected as the safety standard, and assuming a correspondingly larger cross-sectional area, with other factors remaining the same as chosen by the RTF, the registration mass cutoff would have been more than 6 kilograms.)

Another interesting comparison can be made to "non-lethal" (aka "less-lethal") munitions used by police forces worldwide designed to be effective with a low probability of lethality. The table below shows some of the types of non-lethal munitions commonly used, their projectile mass and velocity, the resulting kinetic energy, and the expected probability of fatality using the curve in the MITRE report.

Munition	mass (g)	velocity (m/s)	KE (J)	PoF (%)
12 ga Bean bag	40	90	162	80
HK L104A1/L21A1 PBR	98	72	254	95
ARWEN AR-1	80	74	219	90

It is clear that if the MITRE PoF curve relied upon by the RTF were accurate, all of these non-lethal munitions would actually be expected to have quite a high probability of causing death. But they do not. The U.S. Department of Justice has compiled statistics on the usage of these munitions. In 2004 they published a report on incidents involving these non-lethal munitions over the period 1985 to 2000.[36] This report analyzed 373 separate incidents, involving almost 1000 projectiles fired, resulting in 782 injuries and 8 deaths. More than 80% of the actual

---

[35] See for example ballistic charts for the 30-06 Springfield, the cartridge used by the US Army's M1 Garand rifle. 3315 ft-lbs muzzle energy, 1855 ft-lbs energy at 300 yds. http://guide.sportsmansguide.com/ballistic-chart/federal_charts/30-06Spring762x63mm.html

[36] Hubbs and Klinger, *Impact Munitions Database of Use and Effects*, National Institute of Justice, DOJ, 2004.

injuries were categorized as bruises or abrasions that required no further treatment. The actual fatality rate corresponds to a PoF of 1% when comparing the deaths to the number of injuries. The authors note that in their conclusion that this is in fact an upper limit, because not all uses of impact munitions are reported. "With just eight deaths attributable to (actual) impact munitions in 372 cases where at least one projectile found its intended mark, it is clear that impact munitions rarely produce fatal injuries. As noted above, the current data includes all known deaths in North America caused directly by impact munitions strikes as of May 2000, but nowhere near the entire population of cases where officers shot citizens with impact munitions up to that date. Consequently, the percentage of cases where citizens struck by impact munitions die is substantially lower than the 2.2% figure yielded from the present data. In sum, the likelihood of death from being shot by impact munitions is extremely low…"[37]

In contrast, the model used by the Registration Task Force would estimate that such munitions would kill nearly everyone struck by them, when in fact the fatality rate is only 1% per impact.

Both of these lines of real-world evidence lead us to the conclusion that the PoF curve used by the RTF to estimate the effect of the impact of UAS on persons, and ultimately to calculate a mass threshold of 250g, was completely unsuitable for the analysis, and far too conservative to inform the policy goal.

The Lowest-Risk UAS

Many attempts have been made to define a "lowest-risk" or "harmless" UAS but few agree. Using the above methods, the RTF decided on a threshold energy of 80 J and, with

---

[37] Hubbs and Klinger, page 22

additional assumptions, arrived at an aircraft mass of 250 g. As shown in this paper, one can use the RTF method, with adjustment of a couple real-world reasonable vehicle parameters, to show that 550g, or even 1100g with consideration of impact dynamics, is a more sensible limit.

However, as we have shown, the RTF method starts with an unrealistic estimate of PoF. The MITRE PoF curve is in sharp disagreement with real world data in the form of battlefield data and non-lethal munitions statistics. In particular, the non-lethal munitions data are real life impacts that achieve only a PoF of ~1% at a KE of ~200 J. If one uses the RTF method to calculate a mass limit using a threshold of 200 J, with a drag coefficient and area approximating a larger quadcopter, and the assumption that about one-third of the KE might transfer to the target in an impact, one finds a mass of approximately 2.2 kg would have been the appropriate threshold for UAS registration.

Conclusions

The conclusions that flow from this analysis are twofold. First, the 250g upper threshold selected by the Registration Task Force should be viewed as very conservative given the general policy goals set by that committee concerning the probability of fatality upon human impact. The calculations were based on assumptions dating back 48 years about the lack of medical care in a thermonuclear war. These assumptions, which should have been rejected long ago as a basis for measuring UAS impact risk, result in a far higher estimated fatality rate than is realistic in a modern society. Second, adjusting just a few real-world factors, and accounting for actual kinetic energy transfer, compel a conclusion that the upper weight limit for a "lowest-risk" UAS is nearly an order of magnitude greater: around 2.2 kilograms. Over the past year since the Registration Task Force has issued its report, the 250g threshold has been used, or proposed, in other jurisdictions to define regulatory categories, including imposition of operational limitations and restrictions that are much more burdensome than the simple registration scheme the RTF was asked to help implement. This trend represents a further perpetuation of myths ostensibly based on science and testing data, but in fact based on decades-old conjecture relating to munitions, shrapnel, and nuclear bombs that have little or no application to the question at hand.

Given the faulty assumptions that underlie the selection of 250 g, regulators should be hesitant to adopt a 250 g UAS category without conducting their own rigorous safety analysis based on the desired policy goals. This is not to suggest that a mass-based UAS category is inappropriate. On the contrary, it is far easier to measure mass than to measure other UAS performance criteria, and mass thresholds play a role in many aspects of aviation regulation and transportation regulation more broadly. But a review of the RTF's work shows that its selection of 250 grams is far too low, and far too conservative, to be used to create a lowest-risk UAS regulatory category. Based on a similar approach to risk estimation, with adjustments for real world factors, we propose 2.2 kg as the upper threshold of a "lowest-risk" UAS category.