

A Proposal to Reduce Government Overclassification of Information Related to National Security

Herbert Lin*

INTRODUCTION: THE PROBLEM

Today, many reports have concluded that there is too much classification of information, and many former government officials appear to concur.¹ For example, the 9/11 Commission found that overclassification is a threat to national security because it inhibits information sharing within the federal government and between the federal government and state and local agencies.² Donald Rumsfeld noted in 2005 his long-held belief that “as a general rule . . . too much material is classified across the federal government.”³ Complementing government reports and statements from former government officials Elizabeth Goitein and David M. Shapiro found that in response to a public request for particular classified government records to be declassified, the relevant agency found that in 92% of all such cases at least some of the requested records need not remain classified.⁴

And these concerns are still being expressed. As this paper is being written, Representatives Duncan Hunter and Martha Roby are requesting the Government Accountability Office to review the government’s classification systems and to examine the degree to which material is classified even when such material does not impact national security.⁵ Senator Jeanne Shaheen (D-N.H.) is calling on the Obama administration to increase transparency by reducing the number of classified documents to reduce costs and combat “a culture of

* Chief Scientist, National Research Council’s Computer Science and Telecommunications Board, but this institutional affiliation is provided for identification purposes only. This document does not represent the views of the National Academies or any unit within the Academies, and is not a document of the National Academies or the National Research Council. I acknowledge with appreciation the inputs of Susan Landau, Judith Reppy, Jack Goldsmith, and Steven Aftergood on earlier versions of this document. © 2014, Herbert Lin.

1. Classified information is defined as “[government] information [that] reasonably could be expected to result in [identifiable or describable] damage to the national security.” See Exec. Order No. 13,526, 3 C.F.R. 298 (2009).

2. NAT’L COMM’N ON TERRORIST ATTACKS UPON THE U.S., THE 9/11 COMMISSION REPORT: FINAL REPORT OF THE NATIONAL COMMISSION ON TERRORIST ATTACKS UPON THE UNITED STATES (2004).

3. Donald H. Rumsfeld, Op-Ed., *War of the Words*, WALL ST. J., Jul. 18, 2005, at A12. For a more complete history of reports indicating a significant problem with overclassification, see Steven Aftergood, *Reducing Government Secrecy: Finding What Works*, 27 YALE L. & POL’Y REV. 399 (2009) [hereinafter *Aftergood* (2009)].

4. ELIZABETH GOITEIN & DAVID M. SHAPIRO, BRENNAN CTR. FOR JUSTICE AT N.Y.U. SCHOOL OF LAW, REDUCING OVERCLASSIFICATION THROUGH ACCOUNTABILITY 1 (2011), available at http://www.brennancenter.org/sites/default/files/legacy/Justice/LNS/Brennan_Overclassification_Final.pdf.

5. Letter from Duncan Hunter to GAO, June 19, 2013 and letter from Martha Roby to GAO, July 11, 2013, both available at <http://www.fas.org/sgp/gao/hunter-roby.pdf>.

secrecy that is antithetical to our democratic traditions and undermines public confidence in our institutions.”⁶

The Public Interest Declassification Board, established by the implementing memorandum for Executive Order 13,526 (“Classified National Security Information”),⁷ found in its 2012 report that “present practices for classification and declassification of national security information are outmoded, unsustainable and keep too much information from the public.”⁸

Why is this the case? Goitein and Shapiro identify several incentives for excessive classification, including:⁹

- A culture of secrecy in government agencies. In such a culture, classified information is deemed more valuable than unclassified information, and classification is used as a mechanism to protect agency influence.
- Concealment of information that reveals governmental misconduct or incompetence. Although classification cannot be used in order to “conceal violations of law, inefficiency, or administrative error [or] prevent embarrassment to a person, organization, or agency,”¹⁰ it is obvious that government agencies have incentives to classify information for exactly such reasons.
- Facilitation of policy implementation. To the extent that knowledge of government actions can be limited only to those individuals who must take action, such actions can take place with minimal debate or delay. This is especially true when such actions might be controversial if made public.
- Fear of repercussions for failing to protect sensitive information. While individuals are likely to be criticized for such failures, they are unlikely to be criticized for excessive classification.
- Other demands on classifiers’ time and attention. Making the determination that a given piece of information “could *reasonably* be expected”¹¹ to harm national security entails an inherently difficult and time-consuming process. Thus, when working under time pressure and with a large volume of information about which classification decisions must be made, a classifier has incentives to err on the side of classification.

The idea that these factors incentivize classification has the ring of truth to

6. Press Release, Senator Jeanne Shaheen, Shaheen Calls on White House to Increase Transparency, Reduce Wasteful Spending by Reforming Classification System (May 30, 2013), <http://www.shaheen.senate.gov/news>.

7. Classified National Security Information, Exec. Order No. 13,526, 75 Fed. Reg. 707 (2009).

8. PUB. INTEREST DECLASSIFICATION BOARD, TRANSFORMING THE SECURITY CLASSIFICATION SYSTEM 1 (2012), *available at* <http://www.archives.gov/declassification/pidb/recommendations/transforming-classification.pdf>.

9. GOITEIN & SHAPIRO, *supra* note 4, at 21-27.

10. Exec. Order No. 13,526 § 1.7(a)(1), 75 Fed. Reg. at 710 (emphasis added).

11. *Id.* at § 1.4, 75 Fed. Reg. at 708.

many people within the national security community, although a more nuanced view might be that “although overclassification is widespread, it doesn’t happen very much in my office/agency/program, because we’re on top of it, and besides, the topics in my portfolio are *really* important, unlike the topics handled in all of those other offices/agencies/programs.”

The Goitein/Shapiro analysis of incentives omits one major factor driving classification – the fact that classifiers (and their agencies) incur no monetary cost for classifying information. In traditional economic terms, classification is a free good,¹² and according to basic economic theory, free goods are inevitably overused in the absence of countervailing factors. This economic perspective – even by itself – provides a plausible explanation for the overclassification of information.

I. A STEP TOWARD ADDRESSING THE PROBLEM

If one accepts the premise that too much national security information is classified, it is proper to seek ways to reduce the amount of classified information produced and retained.

A. *Some Mechanisms for Reducing the Amount of Classified Information*

Steven Aftergood has argued that government agencies sometimes have compelling interests in certain information. An agency may wish to declassify information because such information will help to establish legitimacy, reassure its supporters and assuage critics, or counter errors in the record.

According to Aftergood, a number of historical examples demonstrate such motivations in action: the declassification of records relating to the Kennedy assassination, to the Roswell crash of 1947, and to the Office of Legal Counsel’s memoranda on the use of coercive interrogation techniques.¹³ Aftergood also cites the Fundamental Classification Policy Review undertaken by the Department of Energy (DoE) in 1995, which resulted in the declassification of the complete list of U.S. nuclear explosive tests and the release of certain information on the history of U.S. production of fissile materials and inertial confinement fusion technology.¹⁴

In addition, Aftergood notes that the existing classification system has a variety of mechanisms that can be – and to varying degrees, are – used to address overclassification.¹⁵ Two of the most important are:

12. The term “good” is used in an economic sense as a benefit that accrues to the party responsible for classifying a document.

13. Steven Aftergood, *An Inquiry into the Dynamics of Government Secrecy*, 48 HARV. C.R.-C.L. L. REV. 511, 523-524 (2013) [hereinafter *Aftergood (2013)*].

14. Aftergood (2009), *supra* note 4, at 409-410.

15. *Id.* at 407.

- Processes mandated by the Freedom of Information Act (FOIA). Members of the public can file requests to an agency to declassify specific documents, and the agency is required to consider such requests. In some fraction of these cases, these requests are at least partially successful. FOIA also requires agencies to establish an appeals process in which petitioners can ask for a reconsideration of unfavorable decisions.
- The Interagency Security Classification Appeals Panel (ISCAP), which was established in 1995 by Executive Order 12958.72 and is responsible for considering appeals by the public when the originating agency denies a request for declassification (change to active voice) (that is, when all in-agency appeals have been exhausted). Data provided by the Information Security Oversight Office indicates that the ISCAP has declassified “all or some information in a clear majority of the disputed cases it reviewed, even though the classifying agency had refused to do so.”¹⁶

Starting from these historical examples of and successful mechanisms for declassification, Aftergood argues that what they have in common is “a multi-layered process that permits the initial classification decision to be reconsidered from perspectives other than that of the original classifier.”¹⁷ But the two examples above are based on reversing decisions to classify information. An approach of much broader scope is to focus on reducing the amount of information that is classified in the first place.

Classification guides govern what may be classified, at least in principle. In 2009, President Obama asked every executive branch agency with classification authority to conduct a comprehensive Fundamental Classification Guidance Review (FCGR), whose purpose was to ensure that “guidance reflects current circumstances as to what information warrants continued classification” and to identify “information that no longer requires classification and can be expedited for declassification.”¹⁸ The result of this review, completed in June 2012, was that of the 3,103 classification guides reviewed, 869 were either cancelled or consolidated.¹⁹ In addition, from FY 2012 to FY 2013 agencies decreased their total number of original classification decisions²⁰ by 20 percent, and reduced the number of original classifying authorities from 2,269 in FY 2013 to 2,362

16. Aftergood (2013), *supra* note 13, at 526.

17. *Id.* at 525.

18. Exec. Order. No. 13,526, § 1.9(a), 75 Fed. Reg. at 712.

19. Steven Aftergood, *Fundamental Review Leads to Some Reductions in Secrecy*, *SECURITY NEWS* (Aug. 30, 2012), http://blogs.fas.org/secrecy/2012/08/fcgr_reductions/.

20. Classified information is either originally classified or derivatively classified. Original classification refers to the creation of new secrets. Only specific officials known as “original classification authorities” can make the classification decisions that create these secrets, otherwise known as classified information. Derivative classification occurs when those previously created secrets are incorporated in new documents or when appropriately cleared individuals apply classification markings derived from source material or as directed by a classification guide. As a factual matter, the volume of

in FY 2012, results attributed by the Information Security Oversight Office (ISOO) in part to the conduct of the FCGR.²¹ Another such review will begin in 2015, providing an additional opportunity, and of course, there is no reason in principle that such reviews could not happen on a regular basis.²²

Goitein and Shapiro propose an accountability-based approach, which focuses on ensuring that derivative classification conforms to original classification policy.²³ This approach calls for capturing the reasoning underlying classification decisions, audits to review a sampling of such decisions, adverse consequences for classifiers and agencies that violate the intent of classification guidelines, training for classifiers to ensure that they are sufficiently knowledgeable to make appropriate decisions, holding derivative classifiers harmless for failing to follow original classification decisions when such decisions are not clearly conveyed, and providing small financial incentives for individuals to challenge classification decisions.²⁴

Other suggestions to reduce classification include a reduction in the number of authorized classifiers, a maximum classification lifetime of five years for certain kinds of information, legislation that clearly defines what may or may not be classified, and replacement of the existing three-level classification system with a single classification level offering two degrees of protection.²⁵ Suggestions for “disciplined compliance with the rules [for classification]” have also been made.²⁶

B. A Classification Cost Metric for Classified Information

To complement the approaches described above, the proposal below seeks to create serious economic incentives to reduce classification. Consider an approach based on two principles:

- Classification should not be a free good, and some classification cost metric (CCM) should be associated with any piece of information that is designated as classified.
- Those who actually make decisions about classification should benefit from reductions in the amount of classified information produced. If implemented properly, this principle provides incentives for classification decisions that balance the value obtained from classification in any

originally classified information (as measured by the number of classification decisions) is hundreds of times smaller than the volume of derivatively classified information.

21. INFO. SEC. OVERSIGHT OFF., NAT'L ARCHIVES AND RECORDS ADMIN., ANNUAL REPORT TO THE PRESIDENT 1 (2013), available at <http://www.archives.gov/isoo/reports/2013-annual-report.pdf> [hereinafter ISOO 2013 REPORT].

22. Aftergood (2013), *supra* note 13, at 528.

23. Derivative classification is described in *supra* note 20.

24. See GOITEIN & SHAPIRO, *supra* note 4, at 33-49.

25. Aftergood (2009), *supra* note 4, at 404-407 (recounting these various reports).

26. COMMISSION TO REVIEW DOD SECURITY POLICY AND PRACTICES, KEEPING THE NATION'S SECRETS 49 (1985), cited in Aftergood (2009), *supra* note 4, at 416.

specific case against some cost associated with such classification. Furthermore, this principle drives the decision making about classifying versus not classifying to the parties in the system that have the day-to-day responsibilities for such action.

The most basic document containing classified information is written as any other document would be written, except that each and every paragraph, section heading, and figure has a specific classification associated with it, and is marked as such. All information within a paragraph is classified at the level of the specific classification, even if only one particular piece of information within the paragraph is actually classified.

The Appendix presents a notional classified document for those who have not seen such a document in real life.

This proposal defines a classification cost metric (CCM) for the production of a classified document as

$$C = \sum_i C_i W_i$$

in which the subscript i runs from 1 to the number of categories of classification present in the document. For example, i would range from 1 to 3 if a document contained only top secret, secret, and confidential information (that is, 3 categories of classified information). C_0 , which is the cost of protecting unclassified information, is by definition zero.

- \sum_i denotes a summation over the different categories.
- C_i is a cost parameter associated with classification level i . In general, C_i should be expected to increase with the level of classification.
- W_i is the number of words in a document that are protected by a given classification, that is, classification i . W_i is determined by identifying all paragraphs with classification i and totaling the number of words in all such paragraphs.

To illustrate, consider the production of a document that contains only 100 paragraphs and no pictures or other graphics. In these 100 paragraphs are:

- 80 paragraphs with the classification of (U) and which collectively contain 4000 words.
- 15 paragraphs with the classification of (S) and which collectively contain 450 words.
- 5 paragraphs with the classification of (TS) and which collectively contain 200 words.
- C_U is zero by definition.
- Assume in this illustration that C_S is 5 and C_{TS} is 50.

Under these assumptions, the CCM score for this document is calculated as

$$C = (0 * 4000) + (5 * 450) + (50 * 200) = 0 + 2,250 + 10,000 = 12,250.$$

Note the following points regarding the assumptions and calculation:

1. As always, the relationship among the cost parameters is $C_U < C_S < C_{TS}$.
2. The value of C_{TS} is 10 times the value of C_S ; this reflects a judgment that information classified as top secret is 10 times as valuable as information classified as secret.
3. The contribution of C_{TS} to the total cost parameter is larger than that of C_S , even though the volume of secret information is much greater than the volume of top secret information. (Of course, depending on the precise numbers chosen, it would be possible for the CCM contribution of secret information to be greater than the CCM contribution of top secret information for any given document.)
4. Unclassified information contributes nothing to the CCM score.
5. Although the CCM is based on the classification level of each paragraph, the total number of paragraphs at each classification level is irrelevant.

C. Using the Classification Cost Metric in Practice

Scoring a document according to its classification cost metric provides decision makers with a way of judging the relative importance of different classified documents – higher cost means the document is more important and thus improper disclosure would be more consequential. But using the CCM as the basis for limiting classification requires treating the CCM as a kind of currency.

At least three possible uses might be considered:

- A CCM threshold could be established to set an upper bound on the volume of classified information that an office or an agency could produce. That is, the CCM score for every document produced in a given time period could be totaled and the total compared to some threshold. Exceeding that threshold could trigger review of individual classification decisions (such as those described in the Goitein/Shapiro proposal), and since the threshold is arbitrary, it could be raised at will. However, the fact that the threshold can be arbitrarily raised largely defeats any value that the CCM has in forcing decision makers to make choices about classifying versus not classifying information or about using classified information when unclassified information would suffice for the purposes at hand. This use will not be discussed further in this paper.

- The CCM could be used to establish a dollar value for each document, and thus an actual budget associated with the production of classified information can be created and enforced. The use of a dollarized CCM is the focus of the discussion below.
- The CCM can be used to drive decisions about declassifying older documents. An office or agency could earn CCM points towards classifying new documents by declassifying older documents (at a suitable discount rate, explained below). A new classified document could only be issued if sufficient points had been accumulated from declassification of old documents. This use is also described in more detail below.

1. Constraining the Production of Classified Information by Dollarizing the CCM

Dollarizing the CCM requires interpreting C_C , C_S , and C_{TS} as the dollar cost per word of the relevant kind of information. Under this interpretation, the CCM becomes a dollar cost associated with a classified document. (This interpretation means that the specific values for C_S and C_{TS} are no longer notional. Some approach is needed to set their values, and one such approach is discussed later in this paper.)

The second step is for government agencies that produce classified information (whether original or derivative) to establish a total budget for the production of classified information as a line item. This budget is in turn divided among the entities and subentities within the agency that produce classified information. Thus, every office's budget contains a line item for this purpose and no funds from this line item may be used for any other purpose.

The third step is to aggregate the dollarized CCM of every document produced, so that production of classified documents in a fiscal year becomes an expense that the entity must cover with its budget allocation for that year. Notice that as the C_i increase, the number of classified documents that can be produced with a given budget shrinks proportionately (assuming that the average number of words per document remains more or less constant).

The last step is to compare the total classification cost of all classified documents in a given fiscal year to the relevant budget allocation. *If at the end of the fiscal year, the total classification cost of all classified documents produced is below the entity's budget allocation, the office is allowed to keep a fraction β of the cost underrun in the next fiscal year for discretionary but office-related purposes. In the simplest case, $\beta = 1$; that is, the office is able to keep the entire cost underrun for the next fiscal year.*

This mechanism is the heart of the proposal to reduce the production of classified material, as it gives the agency a direct financial incentive to reduce the amount of classified information it produces. This incentive is particularly important in the present budget circumstances, where discretionary funding is highly constrained in agencies and the projects they support.

If the total classification cost of all classified documents produced exceeds

the entity's budget allocation, it can be covered by reprogramming other funds from other line items, and in the case of truly exigent circumstances, additional funds can be requested through the usual procedures used (e.g., supplemental funding procedures and so on). As an alternative, the procedure for gaining CCM credits by declassifying older documents could be invoked at this point.

For contractors that produce classified information, the procedure is similar. A contract is established that provides a line item for the production of classified information. Deliverables to the sponsoring government agency containing classified information are assessed for their classification costs. If the total aggregated classification costs of every deliverable produced is below the line item, the contractor gets to keep the difference. If it is above, the contractor covers the overrun out of its own funds.

To establish a classification cost for a given classified document (i.e., to determine plausible values for C_S and C_{TS}), one point of departure is reports from the Information Security Oversight Office. In FY 2013, the U.S. government reported around 80.2 million classification decisions, the vast majority of which were derivative classifications.²⁷ The cost of the classification system for FY 2013 is estimated at \$11.6 billion for government-only activities, including those of the intelligence community.²⁸ Dividing the cost of the classification system by the number of classification decisions yields about \$145 per classification decision.²⁹ Assuming that each classification decision involves one paragraph of about 40 words, this sets the scale of classification cost at around \$3.60 per word.

According to the notional valuation above, top secret information is presumed to be 10 times as valuable as secret information. Assuming that the level of classification involved in the classification decisions above is mostly secret (and not confidential or top secret), the classification cost for top secret information would be \$36 per word (that is, $C_{TS} = \$36/\text{word}$). Classification cost for confidential information would be less than \$3.60 per word, so for the sake of argument, this paper assumes that C_C is \$1/word. By the previous paragraph, $C_S = \$3.60$.

Note that *even if it is not possible to establish an appropriate classification cost for a given classified document, such a cost could be assigned more or less arbitrarily, subject to this guideline: the cost per document should be large enough to be a reasonable counterweight to the classification process but small enough that the overall classification costs for the agency not be a very large part of its budget.* That is, it is not unreasonable to characterize any dollarized

27. ISOO 2013 REPORT, *supra* note 21, at 1 (58,974 original classification decisions in FY 2013, down from 73,477 in FY 2012; 80.1 million derivative classification decisions in FY 2013, down from 95.2 million in FY 2012).

28. ISOO 2013 REPORT, *supra* note 21, at 24-25.

29. As an editorial comment, the cost per item of classified information is strikingly low, which is yet another indicator that the volume of classified information is way too large.

CCM associated with a classified document as a fee or a fine for producing classified information.

2. Trading Off Classification of New Documents Against Declassification of Old Documents

The CCM can also be used to incentivize the declassification of old documents. The principle is that new classified documents can be created only when a suitable number of older classified documents are declassified. A suitable exchange ratio R must be established for this approach to work, and the CCM is the basis for operationalizing the use of this ratio.

If the exchange ratio R is 10, the classification of one new document with a CCM score of 10,000 (using the original C_i – that is, not interpreted as dollars per word) would require the declassification of other older classified documents with a total CCM score of 100,000.

To prevent gaming of the system by declassifying older documents that are near the declassification date, the exchange ratio could be increased for such documents. For example, the scoring formula for each document could be reduced by a weighting factor based on the number of years until declassification occurs. More precisely, the raw score for each document could be multiplied by the ratio of the number of years remaining in the classification period to the total classification period. This number is close to 1 if the document is to be declassified relatively quickly, and much smaller if the document is intended to be declassified in a short period of time. Thus, if only a small fraction of the total “protected time” remains before the document is intended to be declassified, the “credit” received from declassification is much less than the full CCM score of the document.

The declassification date described above is relevant to original classification. Documents that are derivatively classified may combine originally classified information from many different sources, each of which has a different “declassify on” date. However, in such cases, the latest “declassify on” date from the various original sources is the “declassify on” date of the derivatively classified document.

This approach of trading the declassification of older documents for the right to classify new ones can also be used as a hedge against overruns of the classification budget. Specifically, an agency could prepare for potential overruns by making an intensive effort to declassify old documents. The CCM scores for these documents (suitably weighted for declassification periods) could be totaled and put into an account against which future classified documents would be charged (at the appropriate exchange ratio) in the event of budget overruns.

II. A SHORT DETOUR – IS THE OVERCLASSIFICATION PROBLEM ONE OF VOLUME?

The previous section argued that overclassification is a problem in the sense that the volume of classified information is too large. But volume is not in-

herently a problem if classified information is properly classified – the sole criterion for classifying information is that its unauthorized disclosure would harm U.S. national security. However, the existence of other incentives for classification suggests that some classified information is either improperly classified, does not need to remain classified, or is too important from the standpoint of democratic governance to be withheld from the public. In this view, the overclassification problem is not one of volume per se but rather one of minimizing the volume of classified information that is classified for non-national security reasons.

In principle, it would be preferable to minimize directly the amount of information that is classified for non-national security reasons. Mechanisms based on the incentives described above for declassification and on reducing the need for original classification and audits to ensure adherence to classification guidelines focus on precisely this category of information. But all of them suffer from a fundamental weakness: they are limited in scope. Mechanisms for declassification by definition work only on items that are already classified, and thus can only operate on specified documents. Audit-based mechanisms depend on audits, which again means that only individual decisions can be challenged. Mechanisms based on the use of a classification guide only address the domain of information covered in a given guide.

The value of the economic perspective described in Section I is that it provides incentives for *individuals* to refrain from using classified information and to classify information, and thus acts as an “invisible hand” (from Adam Smith fame) on the market for production of classified documents. Furthermore, the proposal assumes that when individuals making decisions about producing classified documents face constraints on their ability to produce such information, they will think more carefully about what should and should not be classified and/or whether classified information should be used at all.

The technology-based calculation of CCM scores for various documents is both a strength and weakness of this approach. It is a strength because technology is needed to automate the process and make it easy for CCM scores to be calculated, and, fortunately, this approach can be easily implemented using the technology of today. On the other hand, the technology-based approach cannot make judgments about what is and is not “legitimately” classified (at least not today, and perhaps never). For those judgments, humans are needed now (and likely always will be needed), and for this reason audit-based approaches, revisions of classification guides, and so on are complementary to (and are not replaced by) the CCM-based approach.

III. A STRAWMAN IMPLEMENTATION

To describe how this approach might work from the ground up, consider how the author (named Robin) of a derivatively classified document might approach her work. What follows below is a notional description.

- Robin writes the document in the way she has always written in the past, affixing classification markings as appropriate and indicating meta-data such as date of document creation, date of automatic declassification, and so on.
- She submits the document to a computerized CCM scoring program, which returns a CCM score and/or dollar value for the document, calculated as described above.
- She also views a CCM dashboard available on her computer, which indicates the total CCM score (and costs) of all registered documents created in her office (or program) since the start of the accounting period (such as a fiscal year).
- If she wishes, she revises the document (perhaps to reduce its CCM score), and resubmits it as many times as she desires.
- When she is satisfied with the document, she registers the document with the CCM scoring program. What emerges is a CCM score, a control number for the document, and a cryptographic signature for the entire package, document and metadata. Using this information, the CCM dashboard is immediately updated.
- After the document is registered, Robin can still make changes to the document, and updated CCM scores and signatures are provided. These updated scores are also reflected on the dashboard.

Because the total production of classified documents is tracked and scored in real time, everyone in the office knows when the budget allocation is running out. In addition, the amounts of classification that individuals produce would be easily associated with them, and office supervisors in charge of managing budgets for the office would be in a position to see how much each individual is doing. If desired, the office supervisor could himself allocate classification budgets to individuals as a mechanism for managing the office (and the behavior of those individuals).

IV. SOME OBVIOUS QUESTIONS

A. *Questions about the Underlying Approach*

1. Doesn't This Approach Distort the Way in Which the Authors of Classified Documents Approach Their Writing?

Of course it does. If it had no impact on their writing, it would be useless. It does incentivize authors to think more carefully about the use of classified information and to segregate classified from unclassified information.

But perhaps more to the point, this approach does not impose a *requirement* on authors to change their writing styles – rather, it only provides incentives for doing so.

2. Does This Approach Inhibit Information Sharing by Providing Incentives for Brevity?

A reviewer of an earlier draft of this paper argued that this approach would create a new tension between classification and information sharing, even within the government and among appropriately cleared analysts, because analysts would feel (economic) pressure not to elaborate verbally on classified matters. Taken to an extreme, this reviewer argued, such conciseness of expression could undermine clarity or operational effectiveness.

Under some circumstances, this may indeed be an outcome. But the argument makes a number of assumptions. First, it assumes that good analysis depends primarily on classified information. Although this is true in some cases, it is far from universally true.

Second, the argument implicitly assumes that longer memos and analytical works are more useful than shorter ones. From the standpoint of a reader pressed for time, concise memos are a plus rather than a minus as they are easier to understand and remember.

Third, concerns about operational effectiveness are concerns that classified operationally useful information may be omitted. But as described in the discussion of Question 4, time-urgent communications under most circumstances need not be counted at all.

Finally, even if the sharing of classified information is indeed inhibited from time to time, this occasional loss must be weighed against the much broader gain that relevant unclassified information can be shared without any restriction at all.

3. How Does This Proposal Change the Way in Which Classified Information Is Protected?

The most far-reaching impact of this proposal is to change the underlying basis on which information is protected. Today, the classification of any particular document is set at the level of the most highly classified information in the document. If a document contains 90 unclassified paragraphs, 5 Confidential paragraphs, 4 Secret paragraphs, and 1 Top Secret paragraph, the document is classified as Top Secret.

Under this proposal, a document's CCM score determines its classification sensitivity. In particular, it allows for the possibility that a large number of items classified as Secret might have a greater significance, if improperly released, than a single item classified as Top Secret.³⁰ Documents with more items of secret information would be more sensitive than documents with fewer items –

30. One could imagine a single mention of a classified intelligence source that had a greater significance (sensitivity) than a hundred classified State Department cables. In such cases, nothing about this proposal prevents an analyst from *increasing* the number of classified words to boost the total CCM of the document.

this outcome makes intuitive sense, but cannot be reflected under the current classification system.

The current classification system for government information related to national security was developed long before the advent of modern information technology. In that era, simplicity of administering a system for classifying information would have been an understandable virtue – performing the tasks and computations envisioned here in this paper manually would have been a significant burden. But the use of modern information technology can reduce that burden by orders of magnitude, and approaches that were previously unfeasible are now practical. In particular, the proposal described above suggests a way to turn the sensitivity of a document into a continuous rather than a discrete variable, and to take into account multiple dimensions of sensitivity.

To move to a more risk-based approach to safeguarding documents, one might establish levels of security for documents based on their CCM scores. For example, documents with CCM scores exceeding 100,000 might have to be stored in GSA-approved containers with supplemental controls; those with scores between 10,000 and 100,000 might require the GSA-approved containers but not the supplemental controls. Those with scores of less than 10,000 might require only non-GSA-approved containers with a built-in combination lock or a non-GSA-approved container secured with a rigid metal lockbar and an agency head-approved padlock.

4. How Does the CCM Approach Affect Classification Needed Under Operationally Pressing Circumstances?

Under the most basic scheme for implementing a CCM-based approach, only the classified documents in existence at a given time of measurement would count against any total. Thus, if a document were classified and then declassified quickly (i.e., on a time scale of weeks or months), it would not count against a total score and thus would not be charged. This addresses some of the concern about operationally pressing circumstances, as the details of many operations (e.g., the operational orders) need not be kept secret for very long after the operations themselves have occurred.

A more computationally complex method to deal with this problem involves the use of the automatic declassification date – specifically, the number of years that a document is deemed to warrant classification protection. The CCM score for any given document could be increased by a factor f that would increase as the classification period increased. The factor f could have values less than 1 for short times, values greater than 1 for long times, and a value of 1 for intermediate times. Such an approach would tend to drive classifiers to specify shorter periods for classification than would otherwise be the case.

B. Questions about the Mechanics

1. What Is the Burden Placed on Producers of Classified Information to Determine the CCM of a Given Document?

The formula is structured so that a simple computer program could evaluate

the CCM. Moreover, all of the necessary metadata about a classified document should be available (e.g., date of creation, date of automatic declassification).

2. How and to What Extent Can This Approach Be Gamed to Falsely Claim Lower-than-appropriate Costs?

If individuals continue to write documents as they always have, this approach provides a basis on which to score documents for the amount of classified information they contain. But what if individuals change the way they write and format documents to “artificially” reduce the CCM score?

In fact, the CCM approach presumes that individuals will try to game the system in exactly such a manner. To reduce the classification cost of a document, they will minimize mentions of classified information. In some cases, this will mean that the amount of classified information included in the document will be reduced. In other cases, it may mean that the writing of the document is changed to concentrate classified information into fewer classified paragraphs. Such concentration can be used to facilitate declassification review. In still other cases, it minimizes the amount of unclassified information that is comingled with classified information in a given paragraph.

In the extreme case, one could imagine someone trying to lump an enormous number of classified paragraphs into one super-large paragraph and be charged only for one classification decision. But because this approach is based on the number of words and not the number of classification decisions per se, the super-large paragraph would count (to first order) as much as the sum of all of the paragraphs that were aggregated.

3. How Might This Proposal Be Implemented?

As a matter of management, it is unwise to adopt a program with broad-ranging effects without some kind of trial or pilot program to shake out flaws in the proposal. Thus, it would be wise to first implement the system in a limited number of offices before adopting it more widely.

A more limited trial would implement this system simply as a scoring mechanism, with no explicit dollar costs. Without the dollar costs, of course, there are no explicit incentives for changing behavior. But with document scores available, an agency or office would have an increased awareness of the potential distribution of costs across documents and document producers.

An even more limited trial might entail nothing more than the collection of data from offices and agencies about the actual volume of classified information they produce. Such information is not easily available today, but would be a logical first step in testing assertions that the amount of classified information produced is excessive.

C. Questions about Budget and Finances

1. How Should the Overall Classification Budget Be Determined?

The discussion above suggests one approach, based on data contained in

reports of the ISOO. The cost of the classification system for FY 2012 is estimated at \$11.7 billion for government-only activities, including those of the intelligence community.³¹ The FY 2012 appropriations for six key agencies that produce classified information (the Departments of Defense, State (and Foreign Operations), Treasury, and Justice; the National Nuclear Security Administration of the Department of Energy; and the national intelligence community) were \$633.3 billion,³² \$53.5 billion,³³ \$12.2 billion,³⁴ \$27.4 billion,³⁵ \$16.8 billion,³⁶ and \$52.6 billion³⁷ respectively, or a total of about \$796 billion. So as a rough ballpark estimate, one could say that somewhere between 1 and 2 percent of the budget of these agencies is (implicitly) devoted to classification. An overall classification budget might then be established by simply mandating that 1-2% of a program's budget be devoted to it. Such a figure meets the guideline established above – it is large enough to be a reasonable counterweight to the classification process but small enough that the overall classification costs for the agency would not be a very large part of its budget.

Another more complex approach would be to ask agencies (or subentities, such as offices) to estimate, for a given year, the fraction of its budget that is incrementally devoted to protecting classified information. Agencies must undertake such an exercise to report their “classification-related costs” for the annual ISOO reports.

Furthermore, under this proposal, the agency has self-interested incentives to estimate this number accurately. If its estimate is too high, too much of its budget is fenced off, and less money will be available for doing the useful work of the office. If the estimate is too low, it reduces its ability to protect sensitive

31. This estimate includes the sum of the \$9.77 billion reported for 41 agencies not within the U.S. intelligence community and approximately 20% more to account for US Intelligence Community costs, as suggested by the ISOO. *See* INFO. SEC. OVERSIGHT OFF., NAT'L ARCHIVES AND RECORDS ADMIN., ANNUAL REPORT TO THE PRESIDENT 24-25 (2012), *available at* <http://www.archives.gov/isoo/reports/2012-annual-cost-report.pdf>. The reason that FY 2012 figures have been used in this section is that FY 2012 is the latest year for which data is available for an “apples-to-apples” comparison between appropriations figures and spending on classification. For purposes of the rough calculation at hand, the use of FY 2013 figures would not change the outcome significantly.

32. PAT TOWELL, CONG. RES. SERV., R41861, DEFENSE: FY2012 BUDGET REQUEST, AUTHORIZATION AND APPROPRIATIONS (2012), *available at* <http://www.fas.org/sgp/crs/natsec/R41861.pdf>.

33. SUSAN B. EPSTEIN, CONG. RES. SERV., R42621, STATE, FOREIGN OPERATIONS, AND RELATED PROGRAMS: FY2013 BUDGET AND APPROPRIATIONS (2012), *available at* <http://www.fas.org/sgp/crs/row/R42621.pdf>.

34. GARRETT HATCH, CONG. RES. SERV., R42730, FINANCIAL SERVICES AND GENERAL GOVERNMENT: FY2013 APPROPRIATIONS (2012), *available at* <http://www.fas.org/sgp/crs/misc/R42730.pdf>.

35. NATHAN JAMES, JENNIFER D. WILLIAMS & JOHN F. SARGENT JR., CONG. RES. SERV., R42440, COMMERCE, JUSTICE, SCIENCE, AND RELATED AGENCIES: FY2013 APPROPRIATIONS (2012), *available at* <http://www.fas.org/sgp/crs/misc/R42440.pdf>.

36. CARL E. BEHRENS, CONG. RES. SERV., R42498, ENERGY AND WATER DEVELOPMENT: FY2013 APPROPRIATIONS (2013), *available at* <http://www.fas.org/sgp/crs/misc/R42498.pdf>.

37. Press Release, Office of the Director of National Intelligence, DNI Releases Budget Figure for FY 2013 Appropriations Requested for the National Intelligence Program (Feb. 13, 2012), <http://www.fas.org/sgp/news/2012/02/nip-2013.pdf>. The indicated figure is the appropriations *request*, which may be different from the actual amount enacted by law.

information that it needs to protect. More importantly, if the estimate is correct, the classification budget is a wash for the agency. And it gives the agency an incentive to reduce the amount of classified information it produces.

D. Questions about Policy and Law

1. What Are the Adjustments That Policymakers Can Make to This Approach?

The adjustable parameters in this model are the values for:

- C_i , the relative cost of a single word of information classified at level i . The ratios of the various C_i indicate the relative values of information classified at different levels. The C_i can also be specified in dollars per word for a dollarized CCM.
- β , the fraction of the security underrun that an office is allowed to keep in the next fiscal year for discretionary purposes. The underrun is defined as the difference (if positive) between an office's security budget and the aggregate classification cost of all classified documents produced by the office. $\beta = 1$ corresponds to the office being allowed to keep the full amount. $\beta < 1$ allows subsequent classification budgets to be reduced.
- R , the ratio of newly created classified information to information that should be declassified before the newly created classified information can be entered into the system.
- The amount allocated for an overall security budget and how that total amount is allocated to offices.

Other document elements can be accommodated. C_i , the cost of a single word of information classified at level i , generalizes to C_{ij} , where i is the level of classification and j is the type of document element (word, graphics, table, and so on). For example, the appropriate measure for graphics might be a cost per byte of an image. Assuming that a one-megabyte image has about the same value of a paragraph of 100 words, C_{image} should be approximately equal to $0.0001 C_{\text{word}}$ for any given level of classification.

A similar approach can be taken for tables. An appropriate measure for tables might be the number of cells (number of columns times number of rows). Assuming that a 5×5 table (25 cells) has about the same value as a paragraph of 100 words, C_{table} should be approximately equal to $4 C_{\text{word}}$.

It is important to note that after all is said and done to set values for these parameters, their values reflect policy judgments in much the same way that fines to discourage undesirable behavior are set.

2. Isn't Original Classification the Root of the Overclassification Problem?

Although derivative classifiers are obliged to follow the classification markings of original classifiers, the former do have the option of not using classified information. In many cases, some thought and reflection will result in a way to convey an idea without using classified specifics. For this reason, derivative

classification is covered under the proposed CCM approach.

On the other hand, the CCM approach can also apply to original classification. Each agency could establish a budget for producing classified information, abiding by the constraints discussed in Section 2.3. However, in this case, budgetary savings achieved by less originally classified activity should revert to the agency with which these original classifying authorities are associated.

3. Why Should the Costs of Classification Be Charged Directly?

Consider the notion of unfunded mandates. Although there are many definitions of an “unfunded mandate,”³⁸ there is broad agreement that the term includes federal requirements for a given party to take certain actions without compensating that party for expenses incurred in taking such actions. Behaving in accordance with federal regulations and statutes with respect to the production, safeguarding, and declassification of classified information without explicit compensation for the costs of such activities surely meets the definition of “unfunded mandate” and but for tradition and history, the entire classification system would be regarded as such.

4. How Does This Change the Role of Congressional Budgetary Oversight?

Congressional overseers of programs generally want to exercise tight control over an agency’s year-to-year spending, in terms of both the number of dollars actually spent and the purposes for which those dollars are spent. If Congress seeks to control (or worse, recover) the money saved by a reduced level of classification in any given year, the agencies will have weaker or fewer incentives to reduce classification. Using the terminology above, a higher degree of congressional control or recovery corresponds in effect to a value of β that is less than one.

Observers of current congressional politics will also note that the regular budget process (that is, one involving yearly appropriations) has not been operative for a number of years. Nevertheless, the basic idea of allowing agencies to retain funds drawn from unspent classification budgets, across appropriation periods, remains sound. Indeed, in a tight and uncertain budgetary environment, the incentives for reducing spending on classification and saving discretionary purposes are even higher.

5. What Are the Legal Problems With This Approach?

There are potential legal and policy problems with this approach, not the least of which is the fact that even the existence of security classifications for national security information is widely regarded as within the sole purview of the Executive branch. It is Executive Order 13,526 and not any statutory pro-

38. See, e.g., ROBERT JAY DILGER & RICHARD S. BETH, CONG. RESEARCH SERV., R40957, UNFUNDED MANDATES REFORM ACT: HISTORY, IMPACT, AND ISSUES (2013), available at <http://www.fas.org/sgp/crs/misc/R40957.pdf>.

vision of law that prescribes “a uniform system for classifying, safeguarding, and declassifying national security information, including information relating to defense against transnational terrorism.”³⁹

Thus, to the extent that this approach makes it more difficult for the Executive branch to classify information, it is likely that the Executive branch will resist its implementation. But if the Executive branch found value in this approach – as it might well do in pursuing its stated goals of openness and transparency – the scoring approach could be implemented with minimal congressional involvement. On the other hand, matters related to dollarizing the CCM and carrying over amounts from one year to the next – the central concept in incentivizing agencies to reduce the amount of classified information they produce – are squarely within the purview of the Congress.

V. DISCUSSION

As suggested in the Introduction, many analysts and commentators believe the current classification system is broken. For the purposes of this paper, the most important flaw in the current system is that the cost of classification is invisible to those who make decisions about classifying information. While both the current system and the CCM-based approach acknowledge the value of classifying information under some circumstances, the CCM-based approach fundamentally improves the current system in three ways.

First, it makes these costs visible to classification decision makers; by dollarizing the relevant metrics, it allows decision makers to make explicit tradeoffs against other goods that also contribute to national security. Second, by introducing a common currency with which to measure the volume of information, it enables policy makers to focus protective measures on the documents that are the most sensitive. Third, by allowing the amount of funds saved by reducing classification to be spent on a discretionary basis, it provides real budgetary incentives for agencies to actually do so.

This paper advances a proposal that is intended as a point of departure for analysts to consider how to fix the broken classification system of today. The reader is urged to focus on the broad outlines of the argument rather than to argue about the specific numbers that actual implementation would necessarily entail. In addition, the paper discusses a number of factors that can be introduced to correct for various features of classification (e.g., length of time for classification protection). Linear relationships are conceptually the easiest to manage, but any mathematical function can be used as long as the general features of the relationship are preserved.

As for previous related work, the author is aware of only one similar study. In 2004, the JASON study group developed a token-based system for measuring risk and determining the harm that might result from improper access to

39. 75 Fed. Reg. at 707.

sensitive information.⁴⁰ Conceptually, the JASON system is similar to the CCM-based approach, and the JASON system could be used to restrain the production of classified information as well. Nevertheless, the write-up contained in the JASON report did not emphasize that particular application.

CONCLUSION

In many areas, there is broad agreement that too much information is classified. But attempts to reduce the volume of classified information have almost entirely been based on exhortation; excessive classification is essentially never penalized. This proposal takes a “carrot”-based approach that offers the agency or office real incentives to reduce the production of classified information.

40. JASON PROGRAM OFFICE, MITRE CORP. HORIZONTAL, INTEGRATION: BROADER ACCESS MODELS FOR REALIZING INFORMATION DOMINANCE (2004), *available at* <http://www.fas.org/irp/agency/dod/jason/classpol.pdf>.

APPENDIX: A NOTIONAL CLASSIFIED DOCUMENT⁴¹*Basic facts about the U.S.S. Enterprise, NCC-1701-D (U)*

(U) The U.S.S. Enterprise NCC-1701-D was a Galaxy-class starship and the flagship of Starfleet. The fifth starship to be named Enterprise, she was commanded by Captain Jean-Luc Picard.

(U) With a total of 42 decks, the Enterprise-D was twice the length and had eight times the interior space of the Constitution-class ships of over a century earlier. She carried a combined crew and passenger load of 1,012.

(S) Defensive systems included 10 phaser banks, 250 photon torpedoes, and a high capacity shield grid; there are some 4,000 power systems in all aboard ship.

(TS) Full acceleration time from reverse, sub-light impulse through nominal top warp speed, warp 9, was 0.03 milliseconds. The ship achieved maximum speed of warp 9.5 maximum when pursued by the entity known as “Q”; at warp 9.3 its engines had passed the red line.

(C) By Stardate 43205, she had already logged tens of thousands of light-years since setting out around Stardate 41153.7 after Jean-Luc Picard formally became her first captain on Stardate 41124.

(U) An average day aboard ship, according to Lieutenant Commander Data, included four birthdays, two personnel transfers, two chess tournaments, a secondary school play, four promotions, the celebration of the Hindu Festival of Lights – and a birth and wedding. Some 13 planets were represented among the ship’s complement as of Stardate 44247.

(TS) During the Borg incursion of 2366, Decks 23-25 were sliced open by the enemy during the final battle over Earth and requiring a refit at Earth Station McKinley which lasted 5-6 weeks. It was far less damage than the kamikaze warp-driven collision Commander Riker had readied as one last recourse, however. During that refit, the starship received a phaser upgrade as well as damage repair and a dilithium chamber articulation frame.

Source for basic description of Enterprise-D: http://www.startrek.com/database_article/enterprise-d

Source for classification markings: Star Fleet Command

Date of creation: Stardate 45134.2

Declassify on Stardate 58367.5

41. *Enterprise D, U.S.S.*, STAR TREK.COM, http://www.startrek.com/database_article/enterprise-d.
