

Modeling Malicious Domain Name Take-down Dynamics: Why eCrime Pays

Jonathan M. Spring

CERT® Directorate; Software Engineering Institute
Carnegie Mellon University
Pittsburgh PA

Abstract—Domain names drive the ubiquitous use of the Internet. Criminals and adversaries also use domain names for their enterprise. Defenders compete to remove or block such malicious domains. This is a complicated space on the Internet to measure comprehensively, as the malicious actors attempt to hide, the defenders do not like to share data or methods, and what data is public is not consistently formatted. This paper derives an ad hoc model of this competition on large, decentralized networks using a modification of Lanchester’s equations for combat. The model is applied to what is known of the current state of malicious domain activity on the Internet. The model aligns with currently published research, and provides a more comprehensive description of possible strategies and limitations based on the general dynamics of the model.

When taken with the economic realities and physical laws to which the Internet is bound, the model demonstrates that the current approach to removing malicious domain names is unsustainable and destined for obsolescence. However, there are technical, policy, and legal modifications to the current approach that would be effective, such as preemptively populating watch lists, limits on a registrant’s registrations, and international co-operation. The results indicate that the defenders should not expect to eliminate or significantly reduce malicious domain name usage without employing new digital tactics and deploying new rules in the physical world.

I. INTRODUCTION

Malicious domain names cause significant trouble on the Internet, and defenders can and should resist their damage. In deciding the best course of action in eliminating malicious domain names from the Internet, a model of the potential success of various approaches would be a powerful tool. Lanchester’s equations model military combat, but can be modified for this purpose. Lanchester’s equations are themselves a modification of the Lotka-Volterra equations, which model predator-prey interaction. Lanchester’s equations have been critical to the modeling of warfare since their introduction in 1916 [1].

The idealized conception of the equations has many assumptions and they have been modified many times

using various assumptions to better accommodate various types of warfare, such as by [2, 3, 4]. Following [3], the basic combat interaction between a red force R and a blue force B over time, the Lanchester equations are:

$$\begin{aligned} dr/dt &= -K_b b(t); r(0) = R_0 \\ db/dt &= -K_r r(t); b(0) = B_0 \end{aligned} \quad (1)$$

Here, K_b and K_r are non-negative scalars representing the effectiveness of the two forces against each other. The number of active blue soldiers (or airplanes, tanks, etc.), b , in the battle will decrease proportional to the number of soldiers with which red, r , is opposing blue. The number of units changes as a function of time over the duration of the battle. The larger initial force will win, if K_b and K_r are equal. However, if R_0 were twice B_0 , blue could try to compensate by being more effective at destroying red, i.e. increase K_b . But to overcome this 2:1 disadvantage in numbers, K_b would have to be 4 times more effective than K_r . That is, the basic Lanchester equations are second order.

In a military conflict in which both unit types have the same destructive effectiveness, the side with a numerical advantage will have significantly fewer total casualties by the end of the conflict. The effect is demonstrated in the following contrived examples. The Spartans, with superior numbers, achieve a lopsided victory just by maintaining parity with the destructive effectiveness of the smaller force (Figure 1). In order to overcome the Spartans superior numbers, the Athenians must be 4 times more effective to battle them to a draw (Figure 2).

There are many modifications of the Lanchester models. For example, the equations can be modified to take in to account heterogeneous forces (riflemen, infantry, tanks, etc. on each side) each with a different success rate against each type of enemy unit. This can be accommodated by constructing matrices analogous to (1) [2].

Lanchester’s basic equations are derived from the Lotka-Volterra model of predator and prey. Lanchester removed elements that accounted for the birth of prey and the natural death of predators. This is sensible,

Figure 1. Contrived example in which the Spartans have double the initial force and same destructiveness.

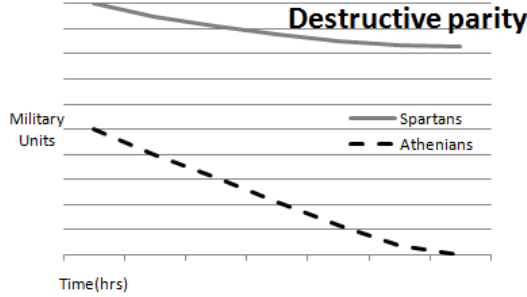
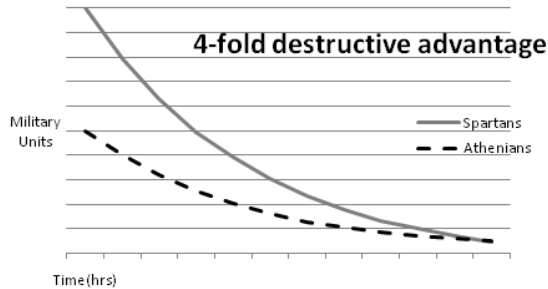


Figure 2. Contrived example in which the Spartans have double initial force but the Athenians have 4-times their destructive effectiveness.



as the time scale of battles does not allow for the production of new soldiers during the battle. The basic predator-prey equations incorporate this concept as follows, where y is the number of predators and x is the number of prey [5]:

$$\begin{aligned} \frac{dx}{dt} &= x(A - By) \\ \frac{dy}{dt} &= -y(C - Dx) \end{aligned} \quad (2)$$

Here, the terms $x(t)$ and $y(t)$ are abbreviated as x and y , respectively. This convention will continue for future variables which are a function of time. The symbols A , B , C , and D are non-negative scalars. The terms A and C account for the birthrate of prey and the natural death rate of predators, respectively. The terms B and D account for the rate at which the predators kill prey and use that energy to reproduce, respectively. Birthrates make (2) conceptually different from (1), as does the concept of consumption for reproduction.

Yet like (1), the Lotka-Volterra equations model entities that both destroy and are destroyed. This is not true of domain names. Therefore equations to model domain dynamics modify both the Lanchester and Lotka-Volterra equations significantly, while still taking their inspiration from the earlier equations.

II. EXISTING EMPIRICAL EVIDENCE

There are observations about the current dynamics of the digital environment and past work that empirically support the design of this model. This evidence is difficult to assemble because there is no agreed-upon reporting format or venue for domain name abuse. Furthermore, data may be considered proprietary or sensitive. However, in the author's experience, the available evidence supports the model as proposed in Section III. Malicious actors can and do automate fixed costs, externalize fixed costs, reduce variable costs, and utilize existent infrastructure. These are generally economically sensible actions [6], and malicious actors can intelligently follow them the same as everyone else [7].

Some of the larger respites from malicious activity have been not due to defensive action, but due to the decisions of the attackers. The group of malicious actors known as Avalanche abandoned their phishing exploits on their own, although presumably to pursue other endeavors [8]. The APWG phishing reports have not measured a large decrease in phishing attacks besides this purposeful abandonment over several years of measurement [9, 10, 8, 11, 12].

Malicious actors' activities can be modeled as having a negative marginal cost and low variable cost due to the dynamics of digital economics. Additionally, many initial costs are borne by other organizations, such as registrars, who sell domain name services. Initial costs may also be borne by previous criminals, who have gone through the trouble to establish botnets usable as name servers or other services. This infrastructure further reduces set up costs, i.e. E in (5). Several large botnets tend to exist at any given time [13].

In practice, the cost of domain names to malicious actors is nearly zero. Dozens of dynamic DNS services provide free domain name registration. Furthermore, many registrars and registries permit "domain tasting," in which a registrant is permitted to use a domain for a few days to get a sense of the traffic available to it before paying. Even though ICANN successfully implemented policies to eliminate this practice within generic TLDs, the country code TLDs aren't bound by the same policies [14]. The useful lifetime of a malicious domain is already below a few days [15]. So the attack patterns are already adapted to making use of a domain well within the time frame afforded by this free domain tasting.

There are other factors that make domains free. Much malicious content is served from hacked servers on hacked IPs or domains. In late 2007 to early 2008, about 75% of phishing domains were hosted on hacked, and re-hacked, servers [16]. Two years later, the rate was independently estimated to be at least

85% of phishing domains hosted on existing domains [17]. The rate of hosting on compromised services is important, because it is nearly impossible to proactively predict what will be hacked. The technical response is therefore pinned in many cases to a reactive approach, however we will demonstrate that a reactive approach has significant short-comings.

Significant efforts have also been made in measuring the economic cost of cybercrime [18]. This is a difficult multi-disciplinary task, and there is still more work to be done. [18] also importantly differentiates between direct and indirect losses to the defenders. The attackers can only monetize a percentage of the direct losses, even though indirect losses are often much larger. However, existing global direct losses are estimated to be on the order of a billion dollars, give or take an order of magnitude. There is also evidence that targeting physical infrastructure, such as banks, can have a significant impact on criminal behavior, regards rogue pharmacies [19].

III. MODELING DOMAIN NAME TAKE-DOWN

The equations introduced can be modified to suitably describe actors competing in the digital world, given the appropriate assumptions are extracted from (1) and (2). For example, the assumption that every ranged unit can target every enemy ranged unit is not physically feasible for a large force. On the Internet, however, this is reasonable. In usual day to day operations, every part of the Internet is supposed to be reachable by every other part of the Internet [20, p. 22]. In this regard, the basic Lanchester equation format should apply to competition interactions that occur over the Internet.

The competition of domain name take-down is quite different from that of armed combat. Notably, the entities being destroyed are not the entities doing the destruction. This exception essentially violates the assumptions that entities in x and y are commensurate. Additionally, the competition is inherently one-sided. This is more similar to a predator-prey relationship than an armed conflict. The malicious domains do not seek to take down benign domains, but to steal information (intellectual property, personal information, credit cards) and resources (money, CPU time for botnets) from the target. For this reason, malicious domain activity or numbers do not directly affect malicious actors' opponents, even though a competition for resources may have indirect effects.

Like prey, which are born and die during the scope of the competition described by the above equations, domain names are definitely registered and taken down during the competition. On the other hand, digital competition contains asymmetries unlike either war or predation. These asymmetries complicate the connected

equations necessary to describe the competition. The following variables are proposed to represent the salient aspects of the competition interaction among domain names, regardless of their particular malicious use.

Variables that are a function of time:

x_δ := number of active malicious domains

x_v := number of malicious domains newly registered during the interval that have not been activated

x_m := resources (either in time, person-hours, or money) malicious actors have available for registering and maintaining domain names

x_c := resources lost by malicious actors due to non-technical socio-political and criminal penalties.

y_m := resources the community or organization makes available to defensive actions, such as taking down or blocking malicious domains.

y_r := resources (time, person-hours, money, intelligence, intellectual property, etc.) lost by the community or organization as a result of fraud, etc., enabled by malicious domains.

Scalars that represent effectiveness or efficiency:

D_\sim := various; convert units of variable to domain units; must be ≥ 0 .

C_\sim := various; converts units of variable to monetary units; must be ≥ 0 .

Values which are modeled as constants:

N := rate at which new domains are registered by malicious actors

E := engineering and development costs.

With these parameters, we can propose several equations following the spirit of the Lanchester and Lotka-Volterra algorithms. Variables in x represent aspects of the attacker, while variables in y represent aspects of the defender.

$$\frac{dx_\delta}{dt} = D_{x_m}x_m + (D_{x_v}x_v) - (D_{y_m}y_m) \quad (3)$$

$$\frac{dx_v}{dt} = -D_{x_v}x_v - (D_{m2}y_m) + N \quad (4)$$

$$\frac{dx_m}{dt} = C_{y2}y_r - (C_{\delta2}x_\delta) - (C_{x_v}x_v) - (C_{x_c}x_c) - E \quad (5)$$

$$\frac{dy_m}{dt} \approx -C_{y1}y_r \quad (6)$$

$$\frac{dy_r}{dt} = C_{\delta1}x_\delta \quad (7)$$

Equation (3) models the rate of change of the population of active malicious domains. The scalar D_{y_m}

represents the effectiveness of take-down measures per unit of resources devoted. It is estimable by observation, in principle. Block listing has been observed to be reasonably effective [21]. The scalar D_{x_m} represents the effectiveness of efforts to maintain active malicious names and their infrastructure. Since not all newly-registered domains are activated right away, some percentage of the registered domains will be activated over time, which is represented by D_{x_v} . Measurements would be necessary to determine this percentage and if it is constant. Studies to this effect are not known, so for the time being it will be approximated as such.

To balance the equations, they must use the same units. The units of x_m and y_m (resources) are not commensurable with those of x_δ and x_v (domains). However, the important aspect of the equation is that efforts to take down the malicious domains are offset by both efforts to keep them alive and the number of new domains that are activated. The units of D_{y_m} , D_{x_m} , etc., could simply be such that they convert units appropriately. This would not change the general dynamics of the equation. Similar conversions will be assumed for all the equations.

Equation (4) models the change in the number of registered but inactive domains. Since x_v represents domains that have been registered, but not used, it is decremented by the number of domains that are activated in (3). For simplicity, domains are modeled to be registered at a rate independent of the other variables by the scalar N . Community take-down efforts could also reduce these domains, but with a different effectiveness coefficient than these efforts effect live malicious domains in (3).

Equation (5) models the resources available to the malicious actors. The scalar C_{y_r} describes a percentage of those resources stolen that can be incorporated into the malicious actors' resources. Scalars C_{δ_2} and C_{x_v} essentially represent the cost of maintaining and registering a domain name, respectively. The scalar E is a cost independent of the number of domain names active. It represents various engineering, setup, human, and organizational costs.

The resources available to the defender must constrain the defensive resources allocated, y_m . Resources allocated to defense are related to resources lost because it is natural to devote more resources to a bigger problem. However, in practice there are many social, political, and economic factors that alter what resources are allocated for network defense, and such non-technical features are not modeled in (6). The lack of non-technical aspects would be most important to the model in (6), so here this modeling choice is most acutely felt.

In principle, the defender's losses may be reduced

by legal action or insurance payments, however these recuperations will not, for the community as a whole, exceed the costs of providing them. For example, any insurance provider will have to charge more for premiums than they give out, or else that provider will become insolvent. Given losses of this nature it is sensible to assume that y_r is monotonically increasing, as in (7).

The change is positive in (7) because it represents increasing losses, rather than decreasing resources. The variable y_r is also presumed to have no limit in (7). So far, fraud losses have not been so great as to overwhelm the economy or resources of whole communities or organizations, but there is certainly some threshold that y_r could reach for which the defrauded entity would cease to be able to function. Such catastrophes are not considered in this model.

The starting resources available to malicious actors are non-zero, since there are certainly some initial resources rolled over from previous crime, digital or not. This starting funding is the value of x_m at $t = 0$. Evaluating the values of the initial conditions is difficult, but for each variable it is greater than zero.

In theory, there are five costs to model in (5). In reality, for international cybercrime at least, C_{x_c} is essentially zero.¹ This is because there are few effective penalties. Therefore, the current realistic model for (5) is:

$$\frac{dx_m}{dt} = C_{y_2}y_r - (C_{\delta_2}x_\delta) - (C_{x_v}x_v) - E \quad (8)$$

From these differential equations relationships between the resources expended by malicious actors and the community can be derived. Malicious actors will gain resources and capabilities, i.e. $\frac{dx_m}{dt}$ will be greater than zero, as long as the following inequality holds:

$$C_{y_2}y_r > C_{\delta_2}x_\delta + C_{x_v}x_v + E \quad (9)$$

That is, if their profits exceed their expenditures. Furthermore, if the number of active malicious domains is to decrease, i.e. $\frac{dx_\delta}{dt}$ in (3) is negative, then the following must hold:

¹There are no international treaties to account for international cybercrime. The bilateral treaties that exist are avoided by the criminals. INTERPOL cannot press charges, and so the lack of international agreement on what constitutes a crime renders the organization ineffective in this arena. The International Criminal Court (ICC) has not been approved by sufficient nations to be considered effective, especially lacking the support of the United States. The main purpose of the ICC is also war crimes, not cyber crimes, and so would require a significant increase in scope before it would be helpful to this particular problem. As such, the term for criminal penalties and costs is effectively zero. Implementing effective international criminal penalties is a necessary long-term solution.

$$D_{y_m}y_m > D_{x_m}x_m + D_{x_v}x_v \quad (10)$$

The costs represented by x_m and C_{x_v} are essentially the variable costs of a domain name to the malicious actor. Variable costs are opposed to fixed costs, which are initial investments. Total costs are variable costs plus fixed costs. Marginal cost is the change in total cost per one more unit of output, i.e. one more domain [6, p 84ff]. In an unconstrained digital economy, reproduction is reduced to copying patterns of bits, which has a variable cost of essentially zero. As more units are produced, initial fixed costs are averaged out over more units produced and so marginal cost is negative and production costs asymptote towards the cost of a new unit. Therefore, unopposed, domains will approach being free to the malicious actor.

IV. OBSERVED CURRENT DYNAMICS

The equations proposed are difficult to measure. This difficulty is not purely technical, but also due to interests protected by various actors and defenders. Sharing information is still difficult. Therefore, the current proposal is based on a large corpus of anecdotal evidence and reasoning, in addition to the above empirical support. These data points suggest the following constraints on the proposed equations. The proposed constraints lead to useful conclusions in understanding how the actions available to the network defense community are circumscribed, and the following section discusses how those constraints can be overcome to improve network defense.

Both take down resistance and activation of domains are actions that can be automated. Automation increases D_{x_m} and D_{x_v} , which is not beneficial for the defender, considering (10). Automation converts these activities into variable costs, rather than fixed costs. The engineering to automate the operations is the fixed cost. Attackers reduce E in (5) in this way, thereby increasing their profits. Digital, automated costs to the malicious actor will approach the cost of copying bits, i.e. zero, unless non-digital costs are imposed by the defenders. In a purely digital competition, the number of domains available to the malicious actor should always exceed the defender's ability to take them down; the variable cost of detecting domains is not zero. Increasing digital costs is an important and necessary defense, but it is not a sufficient defense.

One might expect there to be switching costs involved in new domain names, i.e. that C_{x_v} would incorporate some component of a switching cost for each new domain. However this is not clear. Since the DNS is one ubiquitous protocol designed to minimize

switching costs, the costs are low as long as the malicious code is capable of asking for the correct names. This engineering is non-trivial, but not actually related to the domain names themselves. Further, malicious code has demonstrated the ability to incorporate both updates and outside data [22], and so the process of what names to look up is also able to be automated, thus providing an avenue to significantly reduce cost.

If the malicious actor would actually purchase the domain, the cost of the domains could be charged to fraudulent or stolen credit cards, perhaps those obtained by previous attacks. Even if stolen credit cards are purchased the cost is minimal. The market for credit card credentials is flooded — price is based on the availability of processing time, not on the supply of stolen credentials.² It is reasonable to imagine that some registrars are established by malicious actors for easier dealing in such stolen credentials. Such malicious establishments have been repeatedly observed for other functions on the Internet, such as the infamous Russian Business Network and many others [23]. This behavior also follows a digital information economy. However the fixed costs of establishing such a business are non-digital, and therefore a potentially useful target for defensive actions.

Given this rationale, it is reasonable to estimate values for $C_{\delta 2}$ and C_{x_v} in (5), (8) and (9); the cost of maintaining a DNS structure and registering new domain names each are near zero and getting nearer. The cost of engineering these solutions, E , is also slowly approaching zero as code is reused and existing infrastructure is leveraged. In this purely digital competition, the terms x_δ and x_v effectively drop out of (8). It then simplifies to a depressing expression about the profits of the malicious actors:

$$\frac{dx_m}{dt} \approx C_y 2y_r \quad (11)$$

Equation (11) states that the resources of the malicious actors will only increase in a purely digital competition with the defenders. Additionally, as long as malicious actors control malicious domains y_r should increase as per (9), and so the malicious actors' profits will increase ever faster.

Certain environments may implement defenses, such as a whitelist in the web proxy on what sites employees can visit, that resist the state of affairs that Equation (11) describes. This would increase the cost to malicious actors to attack that organization. However, such policies are not presently feasible on the Internet

²This information is based on anecdotal evidence from an expert, and has yet to be codified. There are probably additional factors influencing the low cost in addition to large supply, such as lack of trust between purveyors.

at large, and it is not clear that such policies are conceptually feasible at such scale. Thus organizations that connect to the Internet must contend with the reality that Equation (11) describes a feature of the present Internet.

V. RAMIFICATIONS

These models make certain statements. From (9) one can surmise that a defensive tactic is to make x_v have to be very high to increase the cost to the malicious actor. A defender could do so by taking down many domains, forcing many new ones to be registered. The take-down rate is (hopefully) increased by community expenditures, i.e. y_m , as scaled by D_{y_m} in (3). Equation (10) permits a simple evaluation of the success of community expenditures based on whether they take down more domains than the malicious actors can maintain and activate.

Equation (10) also demonstrates that community take-down efforts could be resisted if adversaries create domain structures more resilient to take-down or register and activate many new domains. In practice, both tactics are used.³ In light of the dynamics of the Lanchester equations, the defender would seem to have an advantage. If defender resources, y , are much larger than adversary resources, x , then x would have to compensate by a geometric advantage in D_{x_m} (3) and C_{x_v} (5). Therefore defense should be tenable for the defender if D_{x_m} and C_{x_v} are relatively close in value to D_{y_m} (3). In the physical world it is nearly assured that large, technologically-advanced forces will not be grossly out-gunned by ragtag criminals. Yet, the economics of digital information change the landscape significantly.

It is clear from (10) and (11) and the realities in Section II that an approach which attempts to limit criminal activity solely by removing domains used maliciously is ill-fated. Even if the lifetime of a malicious domain were forced towards its cost of production, i.e. zero, any reactive approach cannot actually eliminate the domain before its use. Therefore, reactive block list services alone⁴ cannot reduce the domain lifetime to its cost, no matter how efficient or well-conceived. In order to react, the domain must have been used, which means profit can be generated in that one use on something that was free to obtain. To reduce the domain lifetime to less than one use, a domain-name-based block list must

³For example, the Conficker C virus generated 50,000 new random domains per day, and each infected host would attempt to contact a pseudo-random selection of 500 of these. This randomness is time and effort in algorithm design, i.e. increased x_m , and the high domain volume is a large x_v .

⁴Block lists currently in use (McAfee RBL, Spamhaus, PhishTank, Google safe browsing, etc.) and receiving academic acclaim (EXPOSURE [24], Notos [25], Kopis [26]) are reactive.

be predictive, i.e. remove domains before they are used. This ability to eliminate registered domains before they are activated is modeled in (4) by the term $D_{m_2}y_m$. In the current landscape, D_{m_2} is essentially equal to zero.⁵

A predictive block list could force the average domain lifetime below one use, since some lifetimes would be zero. There are some examples of such a predictive method. One utilizes intelligence from TLD zone files [27], and another uses simply the URL itself without fetching the content [28, 29]. Even though these explore a useful direction, the false positive rate as proposed is too high to be used in production. Some of the URL detection methods have low false positive rates, yet when considered in the context of the base rate fallacy the rates may still be too high [30, ch 9A]. This makes them a useful tool, but prevents existing technology from completely solving this problem.

Defenders should refine such predictive approaches. Reactive blocking fails because the revenue derived from malicious domains, $C_{\delta_1}x_\delta$ (7), will exceed the cost of new domains, $C_{x_v}x_v$ (5). The value derived from activated domains is small, but it only needs to exceed a minuscule cost. A predictive block list would institute a non-zero D_{m_2} scalar in (4), and alter the landscape significantly by preventing some x_v from becoming x_δ . However it is not a complete fix. Compromised domains and services frustrate a predictive list because malicious activity hosted on newly compromised resources cannot be predicted as easily. However all are not completely defeated. Some methods, such as the URL prediction methods mentioned above, would be able to detect URLs on compromised servers, although what degradation, if any, would be present in that case is not explicitly tested in the papers.

Another defensive option would be to implement policies or actions that would increase the switching or registration costs, including time, for domain names. There are currently no realistic barriers to registering tens of thousands of domain names and maintaining them indefinitely, to be activated when necessary. While most second-level domains cost legitimate registrants money, the criminals do not have to use this method. The model reflects this state in that C_{δ_2} , the cost of maintaining domain names in (5), is effectively zero. A poorly resourced attacker may have to actually purchase domains. However an established attacker may use compromised domains and machines. In this case, the resources are stolen and the attacker does not maintain them except to retain illicit access. Different coefficient values could be supplied for different types

⁵In the case of Conficker the community made a concerted effort to block domains before they were used. Conficker variant C eventually overwhelmed these efforts. There has been little effort to work on this in the general case.

of attackers, however this refinement is left as future work.

Investigation of registration activity would have precedent. Banks investigate suspicious withdrawals or deposits of large amounts of cash [31, p. 88-89]. Similarly, registrars and registries would be justified in investigating such anomalous behavior; there is evidence that malicious domains behave differently in the data they handle [32]. Following the example of banking anomaly detection, it is possible to be fruitful without hampering daily users. Simply capping batch registrations at a low number would be a start. However the community would have to decide on a cap as a collective policy decision, because some registrars base their business around processing bulk registrations cheaply and one does not expect that they would willingly give up that business model. This resistance makes a concrete model such as this one a necessary tool for motivating the appropriate policy. Coordination from registrars and registries would make it more difficult to register a domain (C_{x_v}), to maintain it (C_{x_δ}), and increase defenders' ability to take down both active (D_{y_m}) and registered but unused (D_{m2}) names.

In order to increase the terms $C_{x_c}x_c$ and E in Equation 5 (thus decreasing the malicious actors' profit), pressure could be applied in a couple places. One point would be on registrars. Certainly, some registrars do a better or worse job at preventing abuse than others. Establishing or finding an efficient registrar is a kind of engineering cost (E) for the attackers. If such statistics were available to the community, if not publicly, pressure could be applied to those registrars to improve. Offending registrars should eventually lose their authority to register names. Without any such censure process in place, there is no way to prevent rogue registrars from aiding and abetting the criminals.

Other brick-and-mortar institutions that serve to abet cyber criminals should also be sought out and censured or terminated. The criminal process is not an exclusively digital phenomenon, and traditional countermeasures must not be abandoned. For example, services that launder money (such as banks) would be a valuable target. Such institutions have much higher costs, both to replace, switch, and operate, than simple domain names, also increasing the cost of E . Finding, arresting, and incarcerating the criminals would be a deterrent. The models in Section III and the realities in Section II demonstrate that cyber crime cannot be effectively combated by digital means alone. Yet due to ineffective international coordination, current approaches to the problem are almost exclusively digital. This approach is not sustainable. Effective political changes need to occur to make criminal penalties for international cyber crime a reality; that is, increase $C_{x_c}x_c$.

VI. FUTURE WORK

The most important piece of future work would be to conduct computer simulations of the proposed equations and test out various system dynamics. While the above analytic solution is intuitive and agrees with existing anecdotal evidence, simulation would help drive strategic decisions.

Another important aspect to explore is the values for the variables at the initial conditions. Estimates of some may exist, but there is significant room for improvement. Another possible area for exploration could incorporate the interaction of multiple types of forces within each competing entity, as Dolansky does for Lanchester's original equations [2]. Incorporating such heterogeneity may prove useful in the discussion of domain name take down and competition as well. Each take-down technique will almost certainly have varying effectiveness on the various deployments of malicious domains. Additionally, these models may generalize to other types of network behavior. That determination will have to be made for several large classes of network behavior independently.

In regards to (7), while money may be recouped by the defrauded in a legal case, there is still increasing cost to the community in general in order to pay for the police, insurance, and legal activities. This paper does not introduce such players into the model, and such inputs are left to future work. Such additions would likely be an extension of introducing heterogeneity into the model. There is also some point at which a critical point is reached and the defender begins to collapse due to resources lost. To model this a bounded formula for the elements of (7) could be implemented.

Equation (6) is recognized to be incomplete. This equation oversimplifies the rate at which the defender's resources change. While defensive resources are certainly diminished by resources lost to the attacker, there are myriad other influences. These include resources allocated due to social pressure, legal requirements, ideological values, and others. It is also probable that a community's sense of urgency could increase defensive resources related to the number of malicious domains (x_δ), which would further complicate the interaction. Modeling such interactions is left for future work.

Equation (4) is also incomplete. The rate at which new domains are made available to the malicious actor is not constant, so the dynamics of N must also be modeled. The process is probably related to the agent's needs, opportunities, available resources, and other chance factors. The scalar E in (5) is similarly oversimplified, and would benefit from a more thorough analysis.

Identifying actual resources available to malicious actors is another realm of study. This empirical work is

ongoing in certain niches, such as the spam economy [33]. In general, specifying the values of the scalars in the equations and the initial funding and existing malicious domains is future work for empirical study.

Section V recommends identifying and removing rogue registrars. Successfully doing so would not end the competition, but would likely push the battle to dynamic DNS services. Such services provide free domain names, name server access, and services for user-defined subdomains under domains that the dynamic DNS service provider controls. Essentially, they serve as an informal registrar, registry, and name server operator. Dynamic DNS services could be considered rogue registrars which would require censure; however they do not operate under the authority of ICANN or another central authority and so would be more difficult to find and censure.

It also seems that these dynamics also apply to other areas of cyber security. Spam, compromised user accounts, anti-virus signatures, and software vulnerabilities all may fit this dynamical model, although with different scalar values which are probably not nearly zero in the same places. This similarity should be researched, as if this model generalizes to most cyber security problems on the Internet, it would be more useful than just malicious domain names.

VII. CONCLUSION

Digital countermeasures to malicious domains are still necessary. Their effectiveness has been documented. However, as the competition continues, the malicious actors will continue to adapt around digital countermeasures. The models presented demonstrate that malicious actors should be expected to always be able to adapt around digital countermeasures and still profit. Given the necessary features of a digital economy and reactive blocking, the malicious actors will still have revenues exceeding their costs. These digital methods must be accompanied by physical and policy countermeasures to cyber crime and malicious domain name usage. Malicious domains serve as a means to a human economic end. Criminals will operate in the space where they will not be caught or punished. Without effective penalties equivalent to those for traditional crime, one cannot expect cyber crime to cease of its own accord or by digital countermeasures alone.

Acknowledgement

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

This material has been approved for public release and unlimited distribution.

Carnegie Mellon®, CERT® and CERT Coordination Center® are registered marks of Carnegie Mellon University.

DM-0000802

The author would like to thank Soumyo Moitra for his early formative help.

REFERENCES

- [1] C. A. Fowler, "Asymmetric warfare: A primer," *IEEE Spectrum*, March 2006.
- [2] L. Dolanský, "Present state of the lanchester theory of combat," *Operations Research*, pp. 344–358, 1964.
- [3] M. Lauren, "Describing rates of interaction between multiple autonomous entities: An example using combat modelling," tech. rep., New Zealand Defense Force, Defence Technology Agency, September 2001.
- [4] R. Helmbold, "A modification of lanchester's equations," *Operations Research*, pp. 857–859, 1965.
- [5] A. Lotka, *Elements of physical biology*. Williams & Wilkins company, 1925.
- [6] J. Henderson and R. Quandt, *Microeconomic theory: A mathematical approach*. McGraw-Hill New York, third ed., 1980.
- [7] T. Moore, R. Clayton, and R. Anderson, "The economics of online crime," *The Journal of Economic Perspectives*, vol. 23, no. 3, pp. 3–20, 2009.
- [8] R. Rasmussen and G. Aaron, "Global phishing survey: trends and domain name use in 1H2010," tech. rep., Anti-Phishing Working Group, October 2010.
- [9] R. Rasmussen and G. Aaron, "Global phishing survey: trends and domain name use in 1H2009," tech. rep., Anti-Phishing Working Group, 2011.
- [10] R. Rasmussen and G. Aaron, "Global phishing survey: trends and domain name use in 2H2009," tech. rep., Anti-Phishing Working Group, 2010.
- [11] R. Rasmussen and G. Aaron, "Global phishing survey: trends and domain name use in 2H2010," tech. rep., Anti-Phishing Working Group, March 2011.
- [12] R. Rasmussen and G. Aaron, "Global phishing survey: trends and domain name use in 2Q2012," tech. rep., Anti-Phishing Working Group, September 2012.
- [13] Shadowserver, "Statistics – bots and botnets." <http://www.shadowserver.org/wiki/pmwiki.php/Stats/Statistics>, 2013.
- [14] ICANN, "The end of domain tasting: Status report on AGP measures," tech. rep., ICANN, 2009.

- [15] Cyveillance, “The cost of phishing: Understanding the true cost dynamics behind phishing attacks,” tech. rep., Cyveillance, Inc., Arlington, VA, 2008.
- [16] T. Moore and R. Clayton, “Evil searching: Compromise and recompromise of internet hosts for phishing,” *Financial Cryptography and Data Security*, pp. 256–272, 2009.
- [17] J. Spring, “Large scale DNS traffic analysis of malicious Internet activity with a focus on evaluating the response time of blocking phishing sites,” Master’s thesis, University of Pittsburgh, 2010.
- [18] R. Anderson, C. Barton, R. Böhme, R. Clayton, M. van Eeten, M. Levi, T. Moore, and S. Savage, “Measuring the cost of cybercrime,” in *11th Workshop on the Economics of Information Security*, (Berlin), June 26, 2012.
- [19] K. Levchenko, A. Pitsillidis, N. Chachra, B. Enright, M. Félegyházi, C. Grier, T. Halvorson, C. Kanich, C. Kreibich, H. Liu, *et al.*, “Click trajectories: End-to-end analysis of the spam value chain,” in *Security and Privacy (SP), 2011 IEEE Symposium on*, pp. 431–446, IEEE, 2011.
- [20] C. Hall, R. Anderson, R. Clayton, E. Ouzounis, and P. Trimintzios, “Resilience of the Internet interconnection ecosystem,” tech. rep., ENISA, 2011.
- [21] T. Moore and R. Clayton, “Examining the impact of website take-down on phishing,” in *Proceedings of the Anti-Phishing Working Group’s 2nd Annual eCrime Researchers Summit*, pp. 1–13, ACM, 2007.
- [22] E. Rodionov and A. Matrosov, “The evolution of TDL: Conquering x64, revision 1.1,” tech. rep., ESET, Bratislava, Slovakia, June 2011.
- [23] B. Krebs, “Russian business network: Down, but not out,” *Washington Post*, November 7, 2007.
- [24] L. Bilge, E. Kirda, C. Kruegel, and M. Balduzzi, “EXPOSURE: Finding malicious domains using passive DNS analysis,” *Proceedings of the Annual Network and Distributed System Security (NDSS)*, February 2011.
- [25] M. Antonakakis, R. Perdisci, D. Dagon, W. Lee, and N. Feamster, “Building a dynamic reputation system for DNS,” in *19th Usenix Security Symposium*, 2010.
- [26] M. Antonakakis, R. Perdisci, W. Lee, N. Vasiloglou II, and D. Dagon, “Detecting malware domains at the upper DNS hierarchy,” in *20th Usenix Security Symposium*, (San Francisco, CA), 2011.
- [27] M. Felegyhazi, C. Kreibich, and V. Paxson, “On the potential of proactive domain blacklisting,” in *Proceedings of the 3rd USENIX conference on Large-scale exploits and emergent threats*, USENIX Association, 2010.
- [28] J. Ma, L. K. Saul, S. Savage, and G. M. Voelker, “Beyond blacklists: learning to detect malicious web sites from suspicious urls,” in *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 1245–1254, ACM, 2009.
- [29] J. Ma, L. K. Saul, S. Savage, and G. M. Voelker, “Identifying suspicious urls: an application of large-scale online learning,” in *Proceedings of the 26th Annual International Conference on Machine Learning*, pp. 681–688, ACM, 2009.
- [30] W. Stallings, *Network Security Essentials: Applications and Standards*. Prentice Hall, fourth ed., 2011.
- [31] S. Levitt and S. Dubner, *SuperFreakonomics: Global Cooling, Patriotic Prostitutes, and Why Suicide Bombers Should Buy Life Insurance*. Harper Collins, 2010.
- [32] J. Spring, L. Metcalf, and E. Stoner, “Correlating domain registrations and DNS first activity in general and for malware,” in *Securing and Trusting Internet Names 2011*, 2011.
- [33] C. Kanich, N. Weaver, D. McCoy, T. Halvorson, C. Kreibich, K. Levchenko, V. Paxson, G. Voelker, and S. Savage, “Show me the money: Characterizing spam-advertised revenue,” in *20th USENIX Security Symposium*, (San Francisco, CA), 2011.