



The Fermi Paradox and the Aurora Effect: Exo-civilization Settlement, Expansion, and Steady States

Jonathan Carroll-Nellenback¹ , Adam Frank¹, Jason Wright^{2,3,4} , and Caleb Scharf^{4,5}

¹ Department of Physics and Astronomy, University of Rochester, Rochester, NY 14620, USA

² Department of Astronomy and Astrophysics, The Pennsylvania State University, University Park, PA 16802, USA

³ Center for Exoplanets and Habitable Worlds, The Pennsylvania State University, University Park, PA 16802, USA

⁴ PI, NASA Nexus for Exoplanetary Systems Science, USA

⁵ Department of Astronomy, Columbia University, 550 West 120th Street, New York, NY 10027, USA

Received 2019 February 11; revised 2019 July 2; accepted 2019 July 10; published 2019 August 20

Abstract

We model the settlement of the Galaxy by space-faring civilizations in order to address issues related to the Fermi Paradox. We are motivated to explore the problem in a way that avoids assumptions about the agency (i.e., questions of intent and motivation) of any exo-civilization seeking to settle other planetary systems. We begin by considering the speed of an advancing settlement front to determine if the Galaxy can become inhabited with space-faring civilizations on timescales shorter than its age. Our models for the front speed include the directed settlement of nearby settleable systems through the launching of probes with a finite velocity and range. We also include the effect of stellar motions on the long-term behavior of the settlement front which adds a diffusive component to its advance. As part of our model we also consider that only a fraction, f , of planets will have conditions amenable to settlement by the space-faring civilization. The results of these models demonstrate that the Milky Way can be readily filled-in with settled stellar systems under conservative assumptions about interstellar spacecraft velocities and launch rates. We then move on to consider the question of the Galactic steady state achieved in terms of the fraction X of settled planets. We do this by considering the effect of finite settlement civilization lifetimes on the steady states. We find a range of parameters for which $0 < X < 1$, i.e., the Galaxy supports a population of interstellar space-faring civilizations even though some settleable systems are uninhabited. In addition we find that statistical fluctuations can produce local overabundances of settleable worlds. These generate long-lived clusters of settled systems immersed in large regions that remain unsettled. Both results point to ways in which Earth might remain unvisited in the midst of an inhabited galaxy. Finally we consider how our results can be combined with the finite horizon for evidence of previous settlements in Earth’s geologic record. Using our steady-state model we constrain the probabilities for an Earth visit by a settling civilization before a given time horizon. These results break the link between Hart’s famous “Fact A” (no interstellar visitors on Earth now) and the conclusion that humans must, therefore, be the only technological civilization in the Galaxy. Explicitly, our solutions admit situations where our current circumstances are consistent with an otherwise settled, steady-state galaxy.

Unified Astronomy Thesaurus concepts: [Galaxy evolution \(594\)](#); [Milky Way evolution \(1052\)](#); [Astrobiology \(74\)](#)

1. Introduction

The Fermi Paradox has a long history in discussions of the prevalence of alien technological civilizations (i.e., exo-civilizations) in the Galaxy (Webb 2002; Cirkovic 2018). Originating with a lunchtime conversation in 1950 where Enrico Fermi famously asked “where is everybody?” (Eric 1985), the Fermi Paradox was first formalized by Hart (1975) and has since become a standard framework for addressing questions concerning the prevalence of exo-civilizations (but see Gray 2015). Formally the paradox might be expressed as follows: “if technologically advanced exo-civilizations are common, then we should already have evidence of their existence either through direct or indirect means” (Frank 2018). Here we take *indirect* detection to mean the search for technosignatures (Tarter 2007) from distant sources outside the solar systems (e.g., Siemion et al. 2013; Wright et al. 2014), while *direct* detection means material evidence for an exo-civilization’s visit to Earth or our solar system (Davies 2012; Haqq-Misra & Kopparapu 2012; Wright 2018).

Such a distinction between direct and indirect detection is important. In Hart’s formulation of the Fermi Paradox his “Fact A” was the lack of aliens on Earth now. It was Fact A that then led Hart to conclude that no other technological civilizations of any kind exist or have existed. The lack of indirect detection of

exo-civilizations via radio or other signals represents a different constraint on alien life (but see Kuiper & Morris (1977), who use the idea that there should be probes in the solar system as a reason to search for radio communication to those probes from abroad). This apparent absence of signals has been called a “Great” (Brin 1983) or “Eerie” (Davies 2011) Silence. Such silence has been taken by some to serve as its own answer to Fermi’s Paradox (i.e., we do not see them because they do not exist). The assumption in this interpretation of the paradox is that the Search for Extraterrestrial Intelligence (SETI; e.g., Tarter 2001) has been extensive enough to place firm limits on the prevalence of exo-civilizations. This conclusion is, however, unwarranted. Tarter et al. (2010) examined the volume of radio SETI search space, and concluded that only a tiny fraction of the radio SETI parameter space necessary to reach such conclusions has been covered. Wright et al. (2018) amplified this conclusion with a similar calculation and concluded that the situation is equivalent to searching unsuccessfully for dolphins in a small pool’s worth of ocean water and then concluding the ocean was dolphin-free.

The Fact A interpretation of the Fermi Paradox, focusing on their presence on Earth (or at least in the solar system), presents greater difficulties in resolution. One of the first rigorous discussions of the possibility of contact via interstellar probes

was that of Bracewell (1960), and Freitas (1980) and Tipler (1980) extended the idea to include self-replicating probes that would saturate the Galaxy. Hart (1975) had in mind interstellar settlement by the intelligent species itself. Either way, the math is the same: both Tipler and Hart showed that subrelativistic probes sent out by a single interstellar faring species would cross the Galaxy in approximately 6.5×10^5 yr. Given that this is a small fraction of the Galaxy's lifetime, it would seem that a single species intent on visiting or even settling the Galaxy has had ample time to do so. But Freitas (1985) challenged this version of the Fermi Paradox as well, noting that it is not a formal paradox at all, both since it relies on the assumptions that alien life *would certainly* launch such probes, that those probes would be in the solar system today, *and* that we would have noticed them by now.

A number of authors have attempted to explore these nuances. Ashworth (2014) included settlement parameters such as a maximum probe range and a maximum travel speed among his categorization of solutions to the Fermi Paradox. Others have attempted to invalidate or verify the order-of-magnitude results of Hart and Tipler using both general arguments and more detailed models.

Stull (1979) argued that competition among settled systems for the small number of frontier systems would stall the expansion front. Jones (1976) argued that a species might intentionally slow its expansion in the Galaxy by at least an order of magnitude via population restrictions. Jones (1978) expanded on this theme with a Monte Carlo approach, performing the calculation of a settlement wave numerically, under the assumption that the spread of life through the Galaxy is driven by population growth, and that only planets or stellar systems with large populations would spread to other stars. He found rapid progress of the interstellar settlement front with the front moving at 6% of the speed of the ships themselves. Based on these results and the lack of material evidence of ex-civilizations on Earth, Jones conjectured that no interstellar civilizations had yet arisen, consistent with Hart and Tipler. Later calculations by Jones (1981) exploring a wider range of population growth assumptions came to similar conclusions.

On the other hand, Sagan & Newman (1983) argued that the sorts of self-replicating probes imagined by Tipler would be inherently dangerous and uncontrollable, and therefore would not be constructed in the first place. Chyba & Hand (2005) argued that self-replicating probes would be subject to evolution, mutation, and predation much like life is, greatly complicating the analysis, a proposition explored numerically by Wiley (2011).

Regarding the Hart scenario, Newman & Sagan (1981) describe an analytic calculation that reproduced the results of Jones (1978), but found that under reasonable assumptions about low population growth rates, the progression of the wave could be slow and the time to populate the Milky Way would approach or even exceed its age. Roughly speaking, their argument is that to be long-lived, a civilization must have low population growth, but if they have low population growth, they will not settle nearby systems (see also the sustainability solution of Haqq-Misra & Baum 2009).

Sagan & Newman also suggested that the colonization phase of the civilization would necessarily be finite in duration, and found that for durations less than 3×10^6 yr we should not expect Earth to have been colonized. They argued that since longer durations were not plausible, Fact A posed no significant

challenge to the hypothesis that the Milky Way is endemic with space-faring life.

Tipler responded in a series of papers, prompting Sagan & Newman (1983) to present a detailed rebuttal to Tipler and defense of the Newman & Sagan calculation, referring to Tipler's position as the "solipsistic approach to extraterrestrial intelligence." Here they defended their choices of parameters for population growth, including their assertion that only well-populated planets would launch new settlement ships, and that civilizations would have finite colonization lifetimes.

Walters et al. (1980) pointed out that the possibility of life spreading among the stars has implications for the Drake equation (Drake 1965), and computed that for a maximum travel time of 10^3 yr one should multiply N by a factor $C \lesssim 10$ to account for this, effectively arguing that the difficulty of interstellar travel would limit each civilization to no more than 10 additional systems.

Additional work has been performed by Landis (1998), who included a maximum probe range in his percolation model and had some settlements permanently cease sending out probes. Kinouchi (2001) used an analogy to the large number of uninhabited portions of the Earth to derive a simple model for the persistence of uninhabited regions of the Galaxy, a model Galera et al. (2018) refined and expanded.

Björk (2007) and Hair & Hedman (2013) performed numerical calculations in a 3D grid of stars representative of the Galaxy. Cotta & Morales (2009) used a 2D grid of stars to explore a two-stage colonization strategy with fast exploration probes and slower colonization probes. Gros (2005) explored the possibility that the settlement wave would cease for cultural reasons. Zackrisson et al. (2015) explored the settlement patterns of in-progress Galactic settlement to guide observational searches for Type III Kardashev civilizations.

Lin & Loeb (2015) modeled the spread of nontechnological life via panspermia, concluding that spatial correlations among life-bearing exoplanets would provide strong evidence for the hypothesis. Forgan (2009) describes a general numerical model for the rise and spread of life in the Galaxy, suitable for testing a wide variety of Fermi-Paradox-related hypothesis (see references therein). Vukotić & Cirković (2012) expand on Forgan's work with cellular automata theory.

With the exception of Zackrisson et al. (2015), who included both 3D stellar thermal motion and Galactic shear in their calculations, all of these studies assumed that settlement occurs across a static substrate of stars, and most worked in 2D. As pointed out by Brin (1983), Ashworth (2012), and Wright et al. (2014), this assumption is probably fine in the case of rapid settlement of the Galaxy by relativistic probes, but it cannot be used to support any conclusion that there are regions of parameter space in which settlement stalls or large uninhabited regions persist for long times in an otherwise inhabited Galaxy, because it assumes that settlements cannot be spread through the Galaxy by the motions of the stars themselves. This is particularly important in scenarios with short maximum probe lifetimes (meaning probes that are either slow or have short range).

Ashworth (2012) and Wright et al. (2014) also repeated the admonition of Hart (1975) against reaching for solutions to the Fermi Paradox that invoke a permanent lack of interest in settling nearby stars, as done by Newman & Sagan (1981) and Sagan & Newman (1983), which is an example of what Wright et al. (2014) dubbed as the monocultural fallacy. Such solutions

invoke the unknown and unknowable intentions or motivations of exo-civilizations, and so unless a species goes *extinct* we should not suppose that any propensity for colonization should go to zero permanently for all settlements.

Wright et al. (2014) also address the assumption of Newman & Sagan (1981) and Sagan & Newman (1983) that the drive to settle new systems would be driven entirely by population pressure, since interstellar migration can hardly be expected to reduce overcrowding (von Hoerner 1975) and such motivations for settlement probably do not even describe most migrations of humans across the Earth.

In addition to the question of the settlement front speed, one can also ask about the steady-state properties of a galaxy which have already been swept across by civilization-bearing probes. If we assume that civilizations have finite lifetimes, then an eventual balance should be achieved between the settlement of currently empty worlds and the death of civilizations on previously settled worlds. This question bears directly on Hart’s Fact A. If civilizations have a finite duration, then it is possible that Earth was settled some time in the distant past and all traces of that settlement have been erased by geological evolution. In Schmidt & Frank (2019) it was shown that evidence of previous industrial civilizations in Earth’s deep past would likely be found only in isotopic and chemical stratigraphic anomalies and that the experiments needed to pinpoint non-natural origins for such signals had yet to be performed. The question of solar system artifacts of previous civilizations (alien or otherwise) has also been addressed in Davies (2012), Haqq-Misra & Kopparapu (2012), and Wright (2018).

In this paper we focus on the Fermi Paradox in the form of Hart’s Fact A. We first reexamine the issue of settlement front speeds using both analytic and simulation based methods to track the settlement front in a realistic gas of stars.

We next take up the issue of equilibrium models for Galactic habitation. Novel aspects of our study include the explicit inclusion of thermal stellar motions coupled with the possibility that not all worlds are inherently *settleable*. Often there is the assumption that any planet can be terraformed to the specific needs of the settling civilization. But the idea that the purpose of probes is to build habitable settlements and that all stellar systems are viable targets for such settlements goes to the agency of an exo-civilization; in our work we therefore relax this assumption. In addition, some stars may host indigenous forms life, which may preclude settlement for practical or ethical reasons (see, for instance, Kuiper & Morris 1977 who suggested the biology of Earth might be incompatible with that of would-be settling species.) This theme was explored in (spoiler alert) the novel *Aurora* by Kim Stanley Robinson (Robinson 2015) in which even though a world was formally habitable it was not what we would call settleable. Thus we include the possibility that good worlds are hard to find—what we call the *Aurora Effect*—via a pre-settlement fraction, f , in our calculations. Separating the observable space density of stars from the unknown density of settleable represents an important aspect of our study.

Finally we consider the effect of finite lifetimes for civilizations which arise on the settled worlds in our calculations. This allows us to include the possibility that Earth was settled by an exo-civilization at some time in the deep past but no evidence remains of its existence (Schmidt & Frank 2019). By including the temporal horizon over which evidence of such a settlement

would persist, we are able to constrain the Galactic equilibrium fraction.

The plan of the paper is as follows. In Section 2 we introduce and explore an analytic model for the speed of the settlement front including stellar motions and the fraction of settleable systems, f . In Section 3 we present an agent-based numerical model for the evolution of a settlement front. In Section 4 we discuss the results of the numerical model and the implications for the long-term evolution of the settlement of the Galaxy. In Section 5 we discuss an equilibrium model for the Galaxy fraction of settled systems and discuss the resulting implications given evidence implied by geological Earth’s record and in Section 6 we present results of the steady-state numerical models.

2. Dynamic Model

We first consider the speed of a settlement front driven by the spread of intelligent agents (i.e., agents following a set of algorithmic rules) constrained by the limited range of the spacecraft and the dynamical diffusion of the stellar substrate. In this model we assume that expansion proceeds via short-range probes that travel to a nearby system and settle it. A settled system in this model takes on identical properties to the Ur system, launching additional probes to an unsettled nearby system.

For simplicity we fix the maximum speed of settlement probes, v_p (relative to their host systems), as well as the maximum distance a probe can travel, d_p , in the rest frame of the host system. We also model the stellar substrate as having a Maxwellian velocity distribution with an average velocity of v_s , and a mean density of systems, ρ , of which some fraction f are settleable. Probes can be launched from a system with a periodicity, T_p . For each probe launched the intelligent agents target the uninhabited system with the shortest travel time. We also assume for simplicity that systems once settled continue to be so, although later in this study we will consider steady-state models in which settled systems have finite lifetimes.

We can scale the model in terms of the probe range (d_p), velocity (v_p), and travel time ($t_p \equiv d_p/v_p$) to reduce the five parameters described above into three dimensionless quantities:

$$\eta = f\rho d_p^3 \quad (1)$$

$$\nu_s = \frac{v_s}{v_p} \quad (2)$$

$$\tau_p = \frac{T_p}{t_p}, \quad (3)$$

where η is the normalized density of settleable systems and roughly corresponds to the expected number of systems within a range at any given point, ν_s is the relative speed of stellar substrate motions to probe motion and tracks the importance of the velocity of stars in aiding or restricting Galactic settlement, and τ_p is the ratio of the probe launch period to the probe travel time and corresponds to the relative delay due to building probes before they can be launched. For reference, Henry et al. (2018) estimate the density of stars in the solar neighborhood between 0.07pc^{-3} and 0.09pc^{-3} . Assuming $\rho = 0.08\text{pc}^{-3} = 0.0023\text{lyr}^{-3}$ and $v_s = 30\text{km s}^{-1}$, we have $\eta = 2.3f\left(\frac{d_p}{10\text{lyr}}\right)^3$, $\nu_s = 0.01\left(\frac{v_p}{0.01c}\right)^{-1}$, and $\tau_p = 0.1\left(\frac{T_p}{100\text{yr}}\right)\left(\frac{v_p}{0.01c}\right)\left(\frac{d_p}{10\text{lyr}}\right)^{-1}$. We then seek a solution to the motion of the agents expansion front

Table 1
Table of Parameters

Symbol	Definition	Description
f		Fraction of systems that are settleable
ρ		Density of systems
d_p		Probe range
v_p		Probe velocity
v_s		Average velocity of stellar substrate
T_p		Probe launch period (equivalent to the probe assembly time)
t_p	d_p/v_p	Probe travel time
T_c	$(\pi f \rho d_p^2 v_s)^{-1}$	Encounter time between systems due to stellar motions
T_l	$\mathcal{D}_1 T_p + (1 - \mathcal{D}_1) T_c$	Effective probe launch period
T_s		Settlement civilization lifetime
η	$f \rho d_p^3$	Normalized density of settleable systems within probe range
ν_s	v_s/v_p	Velocity of stellar substrate normalized by probe speed
τ_p	T_p/t_p	Probe launch period normalized to probe travel time
τ_c	$(\pi \eta \nu_s)^{-1}$	Encounter time normalized to probe travel time
\mathcal{D}_1	$1 - e^{-\frac{4\pi}{3}\epsilon\eta}$	Odds of another unsettled system being in range and ahead of the settlement front
τ_l	$\mathcal{D}_1 \tau_p + (1 - \mathcal{D}_1) \tau_c$	Expected probe launch period normalized to probe travel time
ϵ	$\frac{1}{4}$	Odds of an upwind system being unsettled (parameter)
ν_l	$1 + 3\tau_l^3 \log \frac{\tau_l}{\tau_l+1} + 3\tau_l^2 - \frac{3}{2}\tau_l$	Average <i>effective</i> probe velocity normalized to probe velocity
\mathcal{V}	2–3.5	Ratio of fastest speed to average speed for the Maxwell–Boltzmann system
ν	$\max[\mathcal{V}\nu_s, \nu_l]$	Front speed normalized to probe speed
$\Delta\xi$	$\nu \tau_l / \ln 2$	Front thickness normalized to probe range

$r(t) = vt$, or in scaled units $\xi(\tau) = \nu \tau$ where $\tau = \frac{t}{t_p}$, $\xi = \frac{r}{d_p}$, and $\nu = \frac{v}{v_p}$. A summary of the parameters of our model is provided in Table 1.

2.1. Approximation of the Expansion Speed

We next derive the expected front speed, ν , as a function of normalized density, η , substrate velocity, ν_s , and launch period, τ_p , by first considering various limiting cases.

2.2. High Stellar Density Limit

We first consider the simplest case where the normalized density of settleable systems $\eta \gg 1$, so that there are always plenty of neighboring systems within range. In the static limit for the stellar substrate, $\nu_s \rightarrow 0$, we would expect expansion to simply occur as fast as the probes can travel (including the time needed to build the next probe). The time to travel the probe range, d_p , and launch another probe would be $\frac{d_p}{v_p} + T_p$. This gives a front speed of $v = \frac{d_p}{d_p/v_p + T_p}$ or in our scaled units $\nu = \frac{1}{1 + \tau_p}$. This assumes that systems are evenly spaced at the probe distance. If we assume that the probe destination is another system randomly located within the sphere of radius d_p ,

the average distance travelled by probes per trip in our scaled units $\xi = \frac{d}{d_p}$ will be the volume averaged radius

$$\langle \xi \rangle = \frac{3}{4\pi} \int d\Omega \int_0^1 \xi^3 d\xi = \frac{3}{4}. \quad (4)$$

The average speed (including the probe launch period) in our scaled units $\nu_p = \frac{v}{v_p}$ will then be

$$\nu_p = \left\langle \frac{\xi}{\xi + \tau_p} \right\rangle = \frac{3}{4\pi} \int d\Omega \int_0^1 \frac{\xi}{\xi + \tau_p} \xi^2 d\xi \quad (5)$$

$$= 1 + 3\tau_p^3 \log \frac{\tau_p}{\tau_p + 1} + 3\tau_p^2 - \frac{3}{2}\tau_p \quad (6)$$

$$\rightarrow \frac{2/3}{2/3 + \tau_p} \text{ for } \tau_p \ll 1. \quad (7)$$

Note the last line means that in the limit $\tau_p \ll 1$ the probe speed is the same as what would be expected for uniform trips of an effective distance of $2/3 d_p$. And when $\tau_p \rightarrow 0$, the front speed is just the probe speed and $\nu_p \rightarrow 1$. However, if probes are slow enough or the time to launch is long enough, the stellar substrate may actually move faster, $\nu_s \gg \nu_p$, in which case the stellar diffusion would control the rate of expansion. When this is the case the fastest stars, in the tail of the distribution, will determine the expansion rate. Combining

these possibilities we can estimate the front speed in the high stellar density limit as $\nu = \max[\nu_s, \nu_p]$.

2.3. Low Stellar Density Limit

At low densities, $\eta \ll 1$, host systems do not typically have other settleable systems in range (d_p) and must wait for the background stellar substrate motions to bring them within range. The frequency of these close encounters can be thought of as collision rates of particles with a radius of $\frac{d_p}{2}$. For particles to collide the total distance between them must be twice their radius—or d_p . This gives a collisional cross section of $\sigma = \pi d_p^2$, and an average encounter time of $T_c = (f\rho v_s \sigma)^{-1}$, or in dimensionless units $\tau_c \equiv \frac{T_c}{t_p} = (\pi\eta\nu_s)^{-1}$. Defining a dimensionless effective probe launch period τ_l we can, therefore, not exceed this encounter period to get $\tau_l = \max[\tau_c, \tau_p]$. Combining this with the high density limit gives a resulting front speed of

$$\nu_l = 1 + 3\tau_l^3 \log \frac{\tau_l}{\tau_l + 1} + 3\tau_l^2 - \frac{3}{2}\tau_l. \quad (8)$$

Also, note that when the density drops below $\eta < \frac{1}{\pi}$, the collision time becomes longer than the stellar drift time $\tau_c > \nu_s^{-1}$ or equivalently $\frac{d_p}{T_c} < \nu_s$. In this limit, probes, regardless of their speed or assembly time, can only advance a distance d_p per encounter time T_c . In this limit, the maximum effective probe speed $\frac{d_p}{T_c}$ is less than the stellar drift speed ν_s —so probe launches/relaunches cannot outpace the stellar motions. Even a probe launched forward from the fastest moving system at the leading edge of the front will land on a system drifting back toward the front and it will not be able to find another system in range fast enough to further advance the front.

2.4. Static Limit

Before discussing intermediate system spatial densities with a dynamic stellar substrate, it is instructive to first consider the static limit $\nu_s = 0$. In this case the front propagation speed is limited by the effective probe speed (including launch times) $\nu = \nu_p$. However if the density of settleable systems η drops below a critical density η_c , the expansion of the front can be halted due to insufficient connectivity among neighbors.

2.4.1. Critical Density

We first examine the effects of neighbor system connectivity via probes. In other words what are the dynamics of probes hopping from neighbor to neighbor. We begin by considering the 1D equivalent of N systems distributed at random along the unit line (so that $\rho = N$) and ask what is the minimum probe range d_p needed to ensure that no gaps exist that exceed the probe range, thereby halting the settlement front. If $N \gg 1$, the gaps are very weakly correlated and the gap sizes will have a beta distribution of $\mathcal{B}(\alpha = 1, \beta = N - 1)$. The odds that any gap is smaller than d_p is given by the cumulative distribution function (CDF) of the beta function which is the regularized incomplete beta function $I_{d_p}(\alpha = 1, \beta = N - 1)$. And the odds that no gap is larger than d_p is 1 minus the odds that every gap is smaller than d_p and is given by $1 - I_{d_p}(\alpha = 1, \beta = N - 1)^N$. Figure 1 shows what we call the failure probability, P_f , for the settlement front as a function of the normalized density (1D equivalent) $\eta = \rho d_p = N d_p$. This is defined as the odds of any




Figure 1. Failure probability for a settlement front to completely cross a 1D domain vs. the density of neighbors normalized to the probe range η . As the density increases, the average gap size decreases as does the probability of failure. As the number N of systems grows, the failure probability for a given average gap size increases as there are more opportunities for a gap to exist that exceeds the probe range.

gap size between systems being larger than the probe range. For normalized densities η such that $P(\eta) \sim 1$ the settlement front can cross the domain. We also show the effect on P_t of adding more systems to a domain of constant size.

In 3D, a single gap cannot stall the front, but in the static limit, we can treat the system as a homogeneous Poisson point process with a volumetric rate parameter of $\lambda = \rho$, which gives an occurrence rate for a volume of radius d of $\Lambda = \frac{4}{3}\pi\rho d^3 = \frac{4}{3}\pi\eta\xi^3$ where we define a normalized distance $\xi = \frac{d}{d_p}$. This gives the probability of finding N neighbors within some normalized distance, ξ ,

$$P_N(\xi) = \frac{\Lambda^N}{N!} e^{-\Lambda} = \frac{\left(\frac{4}{3}\pi\eta\xi^3\right)^N}{N!} e^{-\frac{4}{3}\pi\eta\xi^3}. \quad (9)$$

The probability of having one or more neighbors within a distance ξ is given by

$$D_1(\xi) = 1 - P_0(\xi). \quad (10)$$

And the probability of having N or more neighbors within a normalized distance $\xi = \frac{d}{d_p}$ is given by

$$D_N(\xi) = 1 - \sum_{i=0}^{N-1} P_i(\xi). \quad (11)$$

The differential change in the probability of finding N or more neighbors as a function of ξ is equal to the probability of finding the N th nearest neighbor at a distance ξ . Taking the derivative we get the probability of finding the N th nearest neighbor,

$$\mathcal{P}_N(\xi) = \frac{dD_N(\xi)}{d\xi} = -\sum_{i=0}^{N-1} \frac{dP_i(\xi)}{d\xi}. \quad (12)$$

We can use the recurrence relation

$$\frac{dP_i(\xi)}{d\xi} = \begin{cases} 4\pi\eta\xi^2(P_{i-1}(\xi) - P_i(\xi)) & \text{if } i > 0 \\ -4\pi\eta\xi^2P_i(\xi) & \text{if } i = 0 \end{cases} \quad (13)$$

to write the probability of finding the N th nearest neighbor at a distance ξ ,

$$\mathcal{P}_N(\xi) = -\sum_{i=0}^{N-1} \frac{dP_i(\xi)}{d\xi} = 4\pi\eta\xi^2 P_{N-1}(\xi), \quad (14)$$

and then calculate the average distance to the N th nearest neighbor by taking the mean of the distribution

$$l_N = \int_0^\infty \xi \mathcal{P}_N(\xi) d\xi = \int_0^\infty 4\pi\eta\xi^3 P_{N-1}(\xi) d\xi \quad (15)$$

$$= \left(\frac{4\pi\eta}{3}\right)^{-1/3} \frac{\Gamma(N + \frac{1}{3})}{(N-1)!}. \quad (16)$$

We can then determine what value of η is required so that the average distance to the N th nearest neighbor is equal to the probe range ($l_N = 1$),

$$\eta_N = \frac{3}{4\pi} \left(\frac{\Gamma(N + \frac{1}{3})}{(N-1)!} \right)^3. \quad (17)$$

From this relation we find that when $\eta \geq \eta_1 \approx 0.1700$, the average distance to the nearest neighbor is within the probe range. When $\eta \geq \eta_2 \approx 0.4030$ the average distance to the second nearest neighbor will be within the probe range, and when $\eta \geq \eta_3 \approx 0.6399$, the average distance to the third nearest neighbor will be within the probe range.

These relations allow us to understand how the availability of multiple hops leads to a fully connected space of settleable systems. To that end we consider how the typical number of connected systems changes as a function of the normalized density η . To determine this we created a set of random points in a 3D box (with periodic boundary conditions) and then calculated the sizes of each isolated subregion N_i for various values of η . In this the average number of accessible systems (not counting oneself) is given by

$$N_a = \frac{\sum N_i(N_i - 1)}{\sum N_i}. \quad (18)$$

Figure 2 shows the resulting average number of accessible systems as a function of density for $N = 10^1, 10^2, 10^3$, & 10^4 . For $\eta \gtrsim \eta_1$, systems have on average a handful of other accessible systems. Once $\eta \gtrsim \eta_4 \approx 0.8777$, systems are nearly fully connected. This means the number of accessible systems for settlement is only limited by the total number of systems in the box. Thus $\eta > \eta_c \equiv \eta_4 \sim 0.88$ represents a threshold density for settleable systems past which the settlement front should expand freely with the probe speed.

2.5. Intermediate Density

Finally we consider the transition from the low density limit $\eta < \frac{1}{\pi}$ (stellar diffusion limited) to the high density limit $\eta > \eta_4$ (probe speed limited). For intermediate densities, the actual time for probe launches will either be τ_p , if at least one other unsettled system is in range, or it will be collision time τ_c . The odds that there is another unsettled system in range and positioned to advance the front is

$$\mathcal{D}_1 = 1 - e^{-\frac{4\pi}{3}\epsilon\eta}, \quad (19)$$

where ϵ is the product of the fraction of the volume that will assist in advancing the front and the fraction of systems in that




Figure 2. Average number of accessible settleable systems vs. the space density of such systems. The density is normalized to the probe range (η). Note that systems tend to have a single settleable neighbor when the average nearest neighbor distance equals the probe range $\eta = \eta_1$. The space becomes fully connected, meaning a single original civilization can settle all settleable systems, when the average distance to the 4th nearest neighbor equals the probe range $\eta = \eta_4$.

volume that are not already settled. For the system to be upwind of the front, $\epsilon < 1/2$. And for our models, we find reasonable agreement with $\epsilon = \frac{1}{4}$.

Using this we can refine our effective launch time that would advance the settlement front from $\tau_l = \max[\tau_p, \tau_c]$ to

$$\tau_l = \mathcal{D}_1\tau_p + (1 - \mathcal{D}_1)\tau_c. \quad (20)$$

2.6. Stellar Velocity Distributions

One final caveat, is that while v_s is the average stellar speed, the front in the low density limit will be driven by the fastest stars. Given N stars with velocities taken from the Maxwell-Boltzmann distribution, the fastest moving star will be travelling a few times v_s at $\mathcal{V}(N)v_s$, where $\mathcal{V}(N)$ is the expectation value for the maximum of N systems taken from a Maxwell-Boltzmann distribution. This factor scales fairly weakly with N going from 2 to 3.5 for values of N from 10^2 to 10^5 (see Appendix A.1) This then gives our final model for the front speed

$$\nu = \max[\mathcal{V}v_s, v_l] \quad (21)$$

$$\nu_l = 1 + 3\tau_l^3 \log \frac{\tau_l}{\tau_l + 1} + 3\tau_l^2 - \frac{3}{2}\tau_l \quad (22)$$

$$\tau_l = \mathcal{D}_1\tau_p + (1 - \mathcal{D}_1)\tau_c \quad (23)$$

$$\tau_c = \frac{1}{\pi\eta v_s} \quad (24)$$

$$\mathcal{D}_1 = 1 - e^{-\frac{4\pi}{3}\epsilon\eta}. \quad (25)$$

2.7. Front Thickness

After the leading edge of the front passes a given point, the local fraction of settled systems to total systems X in the average rest frame should grow exponentially until it saturates.

The growth rate should be proportional to the frequency of encounters between settled and unsettled systems. The average encounter frequency (normalized to the probe travel time) is just $\frac{1}{\tau_l}$ and the average encounter frequency between one settled and one unsettled system is $\frac{1}{\tau_l}X(1 - X)$. This gives the doubling rate, which corresponds to a continuous growth rate of the fraction of settled systems of $\frac{\ln 2}{\tau_l}X(1 - X)$. The time evolution of the settled fraction can therefore be written as

$$\frac{dX}{d\tau} = \frac{\ln 2}{\tau_l}X(1 - X). \quad (26)$$

This results in a logistic growth with a growth timescale of $\frac{\tau_l}{\ln 2}$. In the frame of the front, this exponential growth is stretched out spatially by the front speed ν . We expect the front to grow from unsettled to fully settled following a logistic curve with a dimensionless width of $\Delta\xi = \frac{\nu\tau_l}{\ln 2}$. Note in the low density limit, we have $\tau_l = (\pi\eta\nu_s)^{-1}$ and the front speed is $\nu = \mathcal{V}_s$. This gives a width of $\Delta\xi = \frac{\nu}{\pi\ln 2}\eta^{-1} \approx \eta^{-1}$.

3. Numerical Model

To test our analytic model we ran a suite of numerical agent based simulations. During each timestep, settled systems check to see if they are ready to launch a probe. Systems that are ready to launch a probe will target the unsettled system with the shortest intercept time subject to the constraint that the distance to intercept (in the systems frame of reference) is less than the probe range and that the time to intercept is less than the probe travel time. If the time to intercept is longer than the probe travel time, a system waits to launch the probe until the time to intercept is less than or equal to the probe travel time. Probes are not preemptively launched at subprobe speeds toward intercept locations—as those locations can be reached with a shorter trip duration by simply waiting to launch a probe at the probe speed provided the probe has not been launched toward another system in the meantime—or enough time has passed to have built another. This also avoids systems preemptively targeting other systems that they would not be able to settle for a long time, allowing those systems to potentially be settled sooner. Once a system is targeted, it will not be targeted by other probes, and will become settled after the probe intercept time.

3.1. Initial Setup

Initially the systems are randomly distributed within a periodic box with velocities taken from a Maxwell–Boltzmann distribution. The initial distribution of settled systems is a Heaviside function, $\mathcal{H}(x_0 - x)$. This creates a gradient in settled systems that causes the front to naturally propagate to the right (+x). To follow the front over many crossing times, the reference frame is also shifted into one moving to the right at the expected front velocity. Settled systems that leave the left boundary are reused and become unsettled as they wrap around and re-enter through the right boundary. Initially the speed of this front is estimated using the analytic prescription, but if the front comes within the probe range of either the left or right boundary, the simulation is stopped, the front speed is reestimated using the result of the simulation, and the process repeats.

One final caveat is that in the low density diffusion limited regime, the system with the fastest velocity in the +x direction, will eventually reach the front. As it crosses the front it will settle additional systems that it encounters leaving a cone shaped wake of settled systems and locally increasing the speed of the front. In a truly infinite plane parallel model, the front would have several fast stars located at different locations along the front causing the front to be somewhat corrugated as the overlapping wakes from the fastest stars intersect. In the simulations, the transverse periodic boundaries effectively limit the corrugation scale and the normal periodic boundary conditions ensure that the fastest moving star eventually emerges from the front, at which point the front speed naturally increases to match the fastest particle's speed. In our setup, we therefore shift the systems so that the system with the fastest velocity in the +x direction starts at x_0 at the leading edge of the front. This position is chosen to be two-thirds of the distance across the box.

The simulation box volume was chosen to contain $N = 10^4$ habitable systems and simulations were performed for various values of $\eta = [10^{-1}, 10^1]$, $\tau_p = [0.1, 1.0]$, and $\nu_s = [10^{-3}, 10^{-1}]$. The volume (in units of d_p^3) is given by $\frac{N}{\eta}$, which using our standard probe range $d_p = 10$ lyr corresponds to volumes ranging from $(100 \text{ lyr})^3$ to $(464 \text{ lyr})^3$ and densities of habitable systems ranging from 0.04 to 4 times the density of stars in the solar neighborhood. The extents in x , y , and z (in units of d_p) were $[w + 15\Delta\xi, w, w]$, where $\Delta\xi = \nu\tau_l$ was the approximated front width which varied from 4.6 to 117 lyr and w was solved for using the volume constraint. For the various parameters, w varied from 80 lyr to 224 lyr and in all cases w was at least 8 probe travel distances as well as at least 10 times the system separation scale $\eta^{-1/3}$, which went from 4.6 lyr to 22 lyr. In addition if we assume background stellar motions of 30 km s^{-1} , the probe speed varied from $10^{-3}c$ to $10^{-1}c$ giving probe travel times that varied from 100 to 10,000 yr and probe assembly times that varied from 1 to 1000 yr. The simulations were each run for the longer of 100 effective launch periods (τ_l) or 10 probe travel times and varied from 1000 yr to 30 Myr and was sufficient for the system to have reached a steady state.

4. Results

The top panel of Figure 3 shows a typical snapshot of the numerical simulation (projected to 2D). The red circles correspond to settled systems while the blue circles are unsettled. Green systems are also unsettled, but have been targeted by a settled system. The lower panel shows the fraction of settled systems projected onto 1D as well as the fit to the logistic curve $X = \left[1 + \exp\left(\frac{\xi' - \xi}{\Delta\xi}\right)\right]^{-1}$, where X is the fraction of systems that are settled, ξ' is the dimensionless position (normalized to d_p in the approximately comoving frame), and ξ and $\Delta\xi$ are fit parameters that indicate the dimensionless position and thickness of the front (in the comoving frame). The change in the dimensionless position ξ in the approximately comoving frame is then used to measure the front velocity. For each run, we then calculate the location and average thickness from $20\tau_l$ to $100\tau_l$ and calculate the average front speed and thickness over this time period.

Figure 4 shows the resulting front speed from the numerical model over a range of values for η for various combinations of




Figure 3. Agent based model simulation of the expansion of a Galactic settlement front after reaching steady state. The top panel shows a projected 2D snapshot of a 3D simulation with the density normalized to the probe range ($\eta = 0.2783$). Red circles correspond to settled systems, green to targeted systems, and blue to unsettled systems. The settlement front is apparent in the transition from red to blue. The bottom panel shows the 1D fraction of settled systems along the direction of front propagation (+x) and the logistic fit $\left[1 + \exp\left(\frac{\xi' - \xi}{\Delta\xi}\right)\right]^{-1}$ used to determine the position (ξ) and width ($\Delta\xi$) of the front.

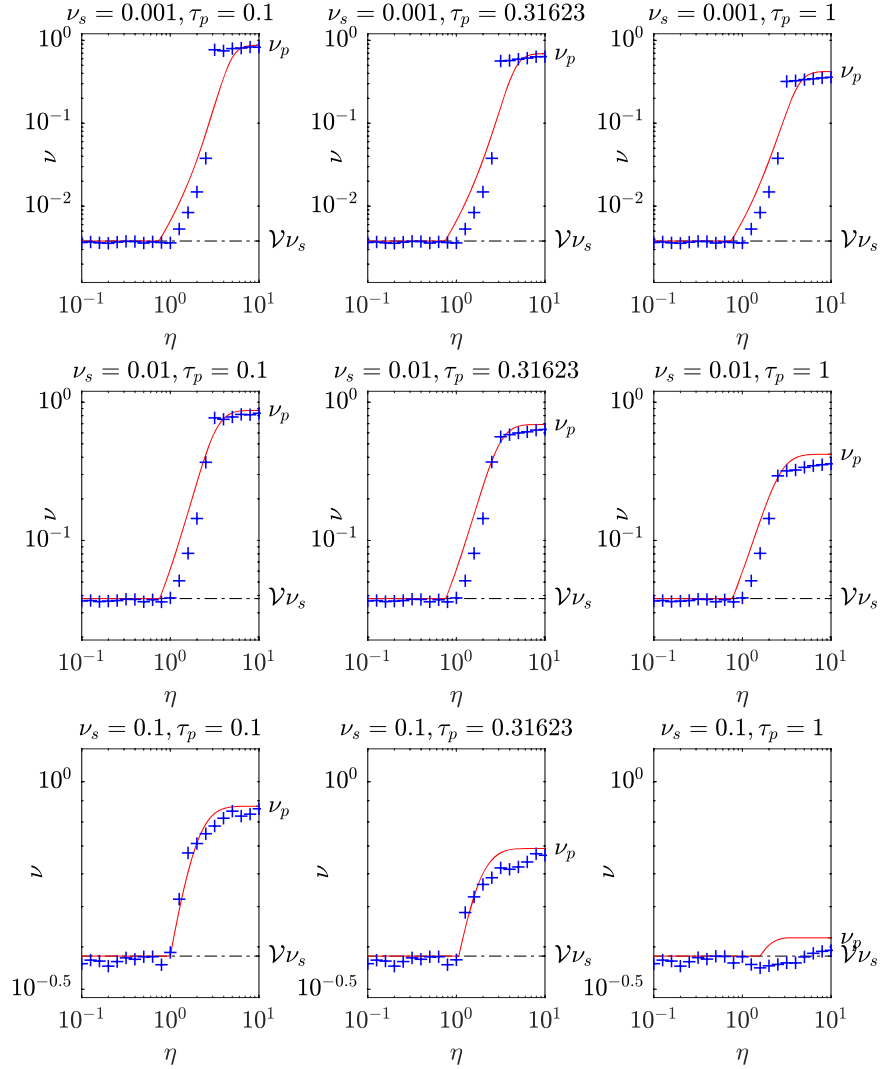


Figure 4. Comparison of analytic and simulation results for the settlement front propagation speed. This figure shows 1D front speeds vs. normalized density. The red line comes from the analytic model and the blue dots come from the simulations. We show runs with the ratio of stellar speeds to the probe speed $\nu_s = [0.001, 0.01, 0.1]$ and the ratios of probe relaunch times to probe travel time $\tau_p = [0.1, 0.31623, 1]$. As η increases, the front speed goes from diffusion limited, to collision limited, to probe speed limited. This figure (and the next) demonstrate the general accuracy of the analytic model discussed in the text.

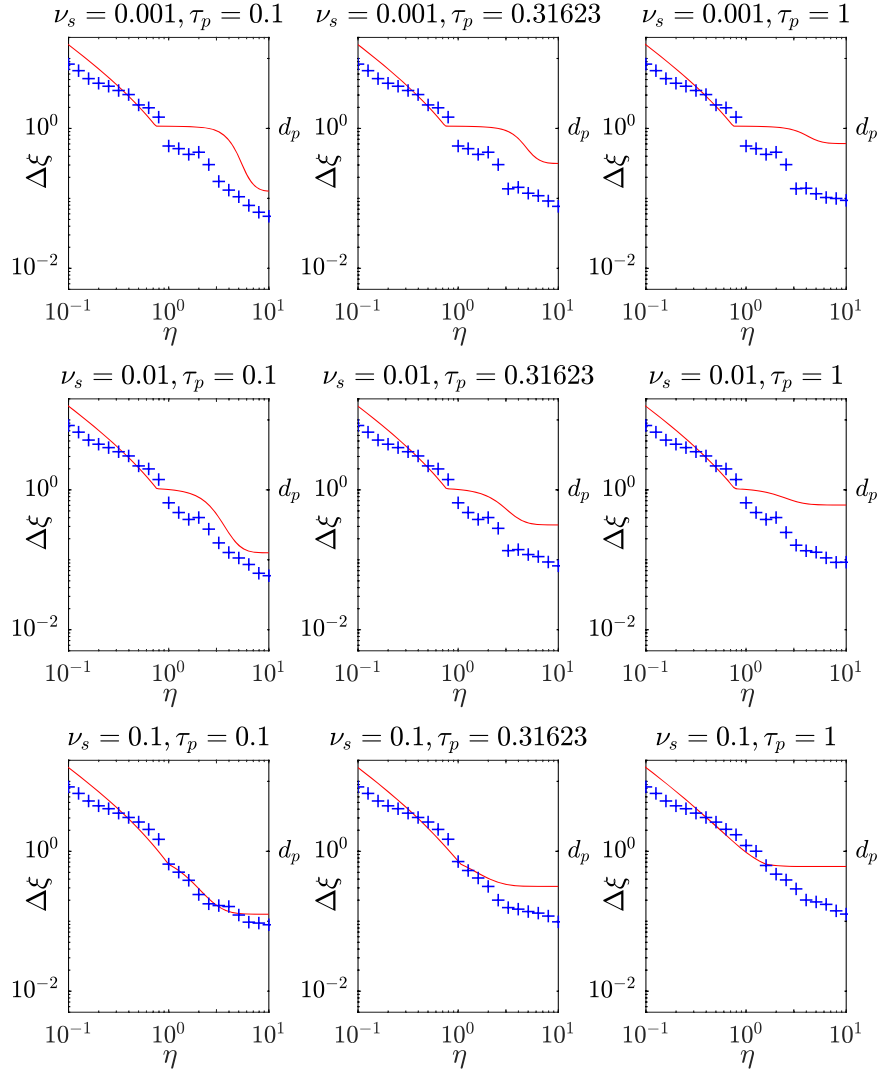


Figure 5. Comparison of analytic and simulation results for the settlement front propagation width. Parameter ranges are the same as in Figure 4. Note that in the low density limit ($\eta < 1$), the launch time is $\tau_l \propto \eta^{-1}\nu_s^{-1}$ while the front speed is $\nu \propto \nu_s$, so the front thickness $\Delta\xi \propto \nu \tau_l \propto \eta^{-1}$ is independent of ν_s .

ν_s and τ_p as well as the analytic estimate. In the low density region, the front speed is just given by $\mathcal{V}\nu_s$. In the intermediate density region, there is a transition from $\mathcal{V}\nu_s$ to ν_p as the effective launch period τ_l transitions from the encounter time τ_c to the probe assembly period τ_p .

Figure 5 shows the measured front thickness as well as our analytic estimate. The analytic estimate does well in the low density regime, where the local growth following the passage of the front may well be described by a logistic curve. However, in the high density regime the front thickness becomes less than d_p and the simple exponential growth model breaks down.

These results demonstrate that the analytic model developed in Section 2 captures most of the important behavior of the settlement front in the low, intermediate, and high settleable system limits. We now use these results to estimate the crossing time of the settlement front across the Galaxy.

4.1. Galactic Crossing Time

We can now apply a very simple order-of-magnitude calculation for the Milky Way, assuming a size of 10^5 lyr and a single stellar velocity dispersion of $v_s = 30 \text{ km s}^{-1}$,

characteristic of the solar neighborhood. While in reality differential rotation, motion of halo stars, and spatial variations in stellar densities and velocities will all be important corrections for realistic models of an expansion of a space-faring civilization, this speed (or rather $\mathcal{V}v_s \approx 100 \text{ km s}^{-1}$) provides a reasonable lower limit for the rate stellar motions can spread life across the Galaxy interior to our Galactic radius. Using this lower limit on speed, we can calculate an upper limit on the Galactic crossing time of 300 Myr. This upper limit is independent of any probe speed, v_p , settleable fraction, f , or probe range, d_p . Figure 6 shows the Galactic crossing time for a range of probe speeds and probe ranges assuming it takes 100 yr to be able to relaunch a probe from a newly settled system. Note in the low density limit $\eta < 1$, this gives 300 Myr. Also as the probe speeds approach the stellar velocities $\nu_s \rightarrow 1$, the front speed becomes comparable to the stellar motions and again the Galactic crossing time goes to 300 Myr. If the probe speed is greater than the stellar speeds ($\nu_s < 1$), and the typical distances to settleable systems is less than the probe range ($\eta > 1$), the Galactic crossing time approaches the light-crossing time 0.1 Myr as the probe velocity approaches the speed of light ($v_p/c \rightarrow 1$). Also note the crossing time tends to increase for shorter probe ranges in




Figure 6. Plot of front crossing time vs. probe speed and settleable system density. The plot assumes a galaxy size of 10^5 lyr, densities similar to the solar neighborhood, stellar speeds of 30 km s^{-1} , and a probe launch period $T_p = 100 \text{ yr}$. Note that the crossing time never exceeds 300 Myr, which is much less than the age of the Galaxy. For reference, *Voyager 1* is traveling at $\sim 10^{-4} c$.




Figure 7. Same as Figure 6, but with the probe launch period increased to $T_p = 1000 \text{ yr}$. Increasing the time between probe launches cannot bring the settlement front crossing time above the age of the Galaxy.

the high density ($\eta > 1$) and high velocity ($v_p \rightarrow c$) limit because there are more frequent hops and the time to relaunch a probe begins to become significant. Figure 7 shows the same but for relaunch periods of 1000 yr.

4.2. Fill-in Time

While the front of settled systems may have had more than enough time to cross the Galaxy, it is worth asking whether the Galaxy has had sufficient time to fill in. Starting from a single technological civilization, what is the timescale for such a civilization to grow to 100 billion civilizations? To settle 10^{11} systems requires only 36 doubling times, so unless the effective probe launch time $T_l = \tau_l t_p$ is greater than $\frac{10 \text{ Gyr}}{\log_2(10^{11})} = 270 \text{ Myr}$, the Galaxy is old enough for every system to have been settled from an initial single civilization. This is of order the crossing time of the entire Galaxy. To restrict launches to once every 270 Myr the encounter time period T_c must be greater than 270 Myr. This requires $(f\rho v_s \pi d_p^2)^{-1} > 270 \text{ Myr}$ or using solar

neighborhood densities and velocities, $\sqrt{f}d_p < 0.071 \text{ lyr}$ or 4520 au. This makes close enough encounters with settleable systems extremely rare.

This result confirms the intuition of Brin (1983), Ashworth (2012), and Wright et al. (2014): using realistic values for stellar motions yields Galactic settlement times shorter than the age of the Milky Way, even for slow ships.

5. Steady-state Model

Given that the Galactic crossing time and potential fill-in time are much less than the age of the Galaxy, we next consider steady-state solutions for a completely settled galaxy. We assume that civilization lifetimes are finite and seek to determine under what conditions settleable systems can be left unsettled for significant periods of time.

If we assume the Galaxy has had time to reach a steady state—and is homogeneous—we can model the ratio of settled to unsettled systems (X) using a simple ordinary differential equation (ODE),

$$\frac{dX}{dt} = \frac{1}{T_l}X(1 - X) - \frac{1}{T_s}X, \quad (27)$$

where T_s is the average lifetime of settlements and T_l is the effective probe launch rate. For our purposes a settlement dies when it ceases to be capable of launching probes. This could be from an extinction event, resource depletion, environmental collapse, or a permanent culture shift to one that does not settle nearby stars.

Note that T_l , the effective probe launch period, is restricted by either the time to assemble a new probe T_p (in the high density limit) or by the time one would have to wait for an encounter with another settleable system within the probe range. As before we set the launch time as the weighted average of the probe assembly time T_p and the collision time T_c ,

$$T_l = \mathcal{D}_1 T_p + (1 - \mathcal{D}_1) T_c \quad (28)$$

$$\mathcal{D}_1 = 1 - e^{-\frac{4\pi}{3}\eta}, \quad (29)$$

where \mathcal{D} represents the odds of a system having at least one neighbor in range. Note we have dropped the factor of $\epsilon = \frac{1}{4}$ since in the steady state we are not concerned with advancing a front, but rather simply launching a probe to any nearby unsettled system.

The factor of $(1 - X)$ represents the odds that an encounter with a settleable system will be with an unsettled system. There is a trivial steady-state solution at $X = 0$ which occurs if civilizations die off before they can launch any probes ($T_s < T_l$). Otherwise, the equilibrium solution occurs at

$$X_{\text{eq}} = 1 - \frac{T_l}{T_s}. \quad (30)$$

In equilibrium, each system must birth (i.e., have an encounter and settle) an average of one unsettled world in their lifetime to compensate for their own death. If systems have several encounters with settleable systems during their lifetime then on average all but one of those will be with other systems that are already settled and the average fraction of systems that are settled will be high. Note that there may be many encounters with systems that are inherently unsetttable, emphasizing our use of the settleable fraction f . If, on the other hand, systems survive just long enough to encounter another




Figure 8. Steady-state fraction of settled systems X vs. settlement civilization lifetime T_s and probe launch times T_l . Note that in the low density limit ($\eta \ll 1$), T_l reduces to the time between encounters with a settleable world due to stellar motions. Dark red regions are fully settled because encounters with a settleable system are frequent (encounter times are less than 1 million years). Dark blue regions are sterile as encounter times are so long that civilizations die before they encounter a system to which they can send a probe. The shaded region corresponds to nonsterile conditions that could be interpreted as being consistent with Earth history (encounter times > 1 Myr; see the text).

settleable system (or rarely two), nearly all of those encounters can be with a unsettled but settleable system.

Figure 8 shows the resulting equilibrium settled fraction as a function of T_l and T_s . In the low density limit ($\eta \ll \eta_c$) the launch time T_l is a function of the settleable fraction f and the probe range d_p (assuming solar neighborhood values for density ρ and stellar substrate velocities v_s). This scaling is shown along the bottom axis. In the figure dark red regions are fully settled because encounters with a settleable system are frequent. Dark blue regions are sterile as encounter times are so long that civilizations die before they encounter a system to which they can send a probe. As we will see in the following sections the shaded region of the plot corresponds to nonsterile conditions that could be interpreted as being consistent with the geologic record of Earth history (see Section 1 of Schmidt & Frank 2019).

6. Steady-state Simulations

To validate the model we ran a sequence of 117 simulations using parameters for the solar neighborhood and explored the dependence on the final settled fraction X on the fraction of settleable systems f and the lifetime for civilizations T_s . We took a constant probe range of 10 lyr and a probe speed of $0.01 c$ and set the probe launch period to $T_p = 0.01 \min(T_s, T_c)$. Each simulation contained $N = 10^4$ particles in a periodic box with a Maxwellian distribution of velocities. We ran each simulation for $20T_c$, which was long enough for the settled fraction X to have reached a steady state. Initially the fraction of settled systems was set to $\max[0.01, X_{eq}]$ randomly distributed throughout the box. This was sparse enough that at early times these initial systems evolved essentially independently. This means one would need ever larger samples to study the

behavior for lower values of the initial settled system population ($X_{eq} \ll 1$).

We then looked at the equilibrium settled fraction (averaged over the final T_c) compared with our analytic model. The right panel of Figure 9 shows the resulting settled fraction. In the high density limit ($\eta > 1$), $T_l = T_p = 0.01T_s$ and we would naively expect $X_{eq} = 0.99$, however the measured settled fractions are lower. This is because in the numerical models, probes are not launched toward systems until after they are no longer settled. And if the travel time is longer than the lifetimes, this delay can create many more targeted systems than settled systems. Appendix A.2 contains a modified model that accounts for this delay and its predictions are shown in the left panel of Figure 9. At both low and high densities it agrees fairly well with the numerical results, though the transition to higher settled fractions appears to be much broader and to have happen much sooner around $\eta \approx 0.2$ as opposed to η_c . This is likely due to a degree of back-filling discussed below.

6.1. Resettlement in the Numerical Models

The numerical models showed two interesting phenomena. First, many of the models with $T_s \ll T_l$ continued to have a handful of systems ($N < 6$) survive for very long times. This was due to pairs of systems very close in phase space with similar velocities and positions that were able to resettlement each other over much longer timescales than their own lifetimes. Looking at the number of unique systems settled over the latter half of the steady-state simulations, we found a clear break between those with only a handful of systems ($N = 6\%$ or 0.06%) and those that used a substantial fraction of the systems (48%). While this back and forth allows for settled systems to exist in the numerical model, we consider any simulation that survives by reusing fewer than six systems as having a settled fraction of zero.

The other interesting phenomena was the excess of systems visible in the right panel of Figure 9 below $T_s = 10^5$ yr for η between 0.1 and 1.0. This too is likely due to an increased probability of resettlement reducing the effective launch period in regions with above average density approaching the critical density. For $\eta = \eta_1 = 0.25$, systems typically have a single neighbor in range allowing for pairs of systems to resettlement one another even if their lifetimes are very short. By $\eta = \eta_2 = 0.4$ there can be subgroups containing ≈ 10 systems that can continually resettlement one another. Within these groups, the probe launch time is just the probe assembly time and if a system dies it can be resettlement by nearby systems. These pockets of settled systems can settle other systems that migrate through. These become the seeds for other pockets to arise if encountered before the system's lifetime. And by $\eta = \eta_4 = 0.88$ systems are fully connected.

Figure 10 shows the results from a run with a settlement lifetime of $T_s = 1.25 \times 10^5$ yr and a encounter time of $T_c = 3.18 \times 10^5$ yr. The density $\eta = 0.1$ gives a neighbor probability of $\mathcal{D}_1 = 0.342$ and an average launch time of $T_l = 2.1 \times 10^5$ yr. While this is still longer than the settlement lifetime, implying $X_{eq} = 0$, a significant number of systems are able to survive in local pockets with higher than average settlement fractions. This is because settled systems (having just been settled) are more likely to have neighbors in range—some of which will have recently become unsettled.




Figure 9. Fraction of settled systems X vs. civilization lifetime T_s and density of settleable systems η . Red corresponds to a fully settled galaxy ($X = 1$) and blue corresponds to a sterile galaxy ($X = 0$). The plot uses solar neighborhood densities and stellar velocities, a probe range of 10 lyr, and velocity of $0.01c$. f is the fraction of settleable systems. The left panel uses the analytic solution while the right is calculated from an array of simulations. The white dot corresponds to the snapshot shown in Figure 10. Note the transition in the settled fraction as a function of density (around $\eta \sim 1$ and $f \sim 1$). For $\eta > 1$ full settlement $X = 1$ might be expected, however short lifetimes can produce low settlement fractions. Note also that the simulations show a higher than expected settlement fraction for densities $\lesssim 1$. This is due to local statistical fluctuations producing overabundances of settleable worlds. These continue to resettle each other after the parent civilization dies.




Figure 10. Settlement simulation snapshot. Blue dots are unsettled systems. The colored circles show settled systems. Systems with the same color share a common ancestor. For these conditions we expect $X \sim 0$ because the effective launch time is longer than the civilization lifetime ($T_l/T_s \sim 2$). Clusters persist, however, because local statistical fluctuations produce overabundances of settleable worlds which can continually resettle one another.

7. Evidence Horizons and Hart’s Fact A

We now turn our attention to Hart’s Fact A which focuses on the question of why Earth has not, apparently, been settled (or at least visited) by another space-faring civilization. The important point to consider is the temporal one. How long ago could Earth have been (temporarily) visited or settled by such a civilization without leaving any obvious trace? If the settlement

occurred 4 billion years ago and lasted for just 10,000 yr would any record of it survive in the geological record?

The answer is: almost certainly not. This implies a temporal horizon over which a settlement might not be seen. We now use the equilibrium solutions to our steady-state model and attempt to constrain the reasoning used in linking Fact A to conclusions about Fermi’s Paradox.

Given assumptions about probe range (10 lyr) and velocity ($0.01c$), we now calculate the typical launch time and equilibrium fraction as a function of the fraction of settleable systems f and the average settlement lifetime T_s (assuming solar neighborhood density of systems and stellar velocities). If encounters between unsettled and settled systems are uncorrelated, we would expect the distribution of times a system remains unsettled (T_u) to follow a Poisson distribution,

$$P(T_u) = e^{-\frac{T_u}{T_l}} X. \quad (31)$$

We can then inquire about the probability of being unsettled for some period that can be compared against evidence in Earth’s geological record. Schmidt & Frank (2019) discuss the mechanisms by which evidence of a previous technological civilizations on Earth would be problematic to find and would likely only exist as potentially ambiguous chemical or isotopic signals, if at all. Motivated by the maximum time span of our simulations, we choose this horizon to be at least 1 Myr. Note that this is a lower limit in that our results are relevant to being unsettled for at least this long. As the horizon time gets longer, the lighter red region in Figure 9 that allows for a reasonable probability of $X > 0$ moves upwards toward longer settled times T_s and to the left toward lower fractions of settleable systems f .

The resulting probability of a given system being unsettled for at least 1 Myr is plotted in Figure 11 versus civilization lifetime T_s and the density of settleable systems η . The left panel shows results for the analytic model; the right panel shows results from the simulations, which validate the analytic model over most of the parameter space we have explored.

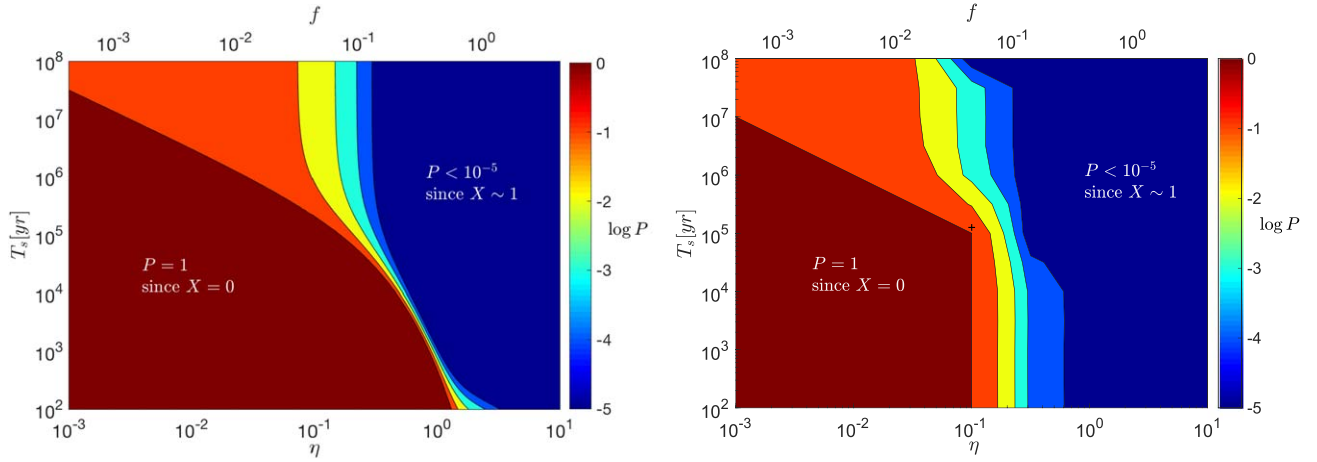


Figure 11. Equilibrium Galactic settlement vs. Earth observation constraints. Here we show the probability of systems being unsettled for at least 1 Myr vs. civilization lifetime T_s and the density of settleable systems η . The left panel shows analytic model results. The right panel shows simulation results. The dark red regions have a probability of 1 because $X = 0$, meaning there are no space-faring civilizations which explains why Earth has not been settled in the last 1 Myr. The absence of civilizations occurs because they die out before being able to encounter another settleable system. The dark blue regions have a very small probability of going 1 Myr without encountering another space-faring civilization. This is because most systems are settled and densities are high enough that encounter times are short.

The dark red regions of Figure 11 show a probability of 1 that the Earth has not been recently visited because $X = 0$ (i.e., because there are no space-faring civilizations, because civilization lifetimes are shorter than the time to encounter another settleable system.) The dark blue regions (indicating a very small probability of Earth having gone 1 Myr or more without encountering another space-faring civilization) represent parameter space where most systems are already settled, and the densities of settleable systems are high enough that encounter times between them and the Earth are short.

The most important point we draw from these plots is that between these two possibilities lies a range of conditions in which the Galaxy supports a population of interstellar civilizations even though Earth would likely not experience settlement attempts for 1 Myr or more. Consider, for example, the situation with a settlement fraction $f = 0.03$ and civilization lifetime of 10^6 . In this case we find it is not terribly unlikely ($P \sim 0.1$) that Earth has remained unvisited for the past 1 Myr or longer, even though $X > 0$. As we will discuss, this result has important conclusions for interpretations of Hart’s Fact A and the Fermi Paradox. Also we note that as the look-back time is extended the analytic version of plots like those shown in Figure 11 tend to have smaller intermediate probabilities for short-lived settlements. We expect however that longer run simulations would retain their wider range in probabilities. We note that the question of evidence horizons should be the focus of its own in-depth future study.

Finally, it is worth pointing out a few features of the probabilities found in our steady-state numerical models (right panel of Figure 11). We found a small probability of Earth remaining unsettled for the past 1 Myr or longer for $T_s < 10^5$ yr for values of $\eta \sim 0.3$ primarily due to small equilibrium values of X supported by resettling near the critical density. This occurs without balancing T_s and T_i . The black “+” in the right panel of Figure 11 represents a run in which a few clusters containing tens of systems tended to persist. In this run the odds of going 1 Myr without a settlement was $\sim 89\%$. Thus if the parameters in this run were to represent the situation in our region of the Galaxy and Earth was not in one of the resettlement clusters, it would be highly probable that we

would not have been settled (or visited) by another civilization for at least 1 Myr. Note also that the narrow transition from $P = 10^{-4}$ to $P = 1$ for $T_s < 10^5$ yr between $\eta = 0.1$ and $\eta = 0.3$ seen in the left panel of Figure 11 could not be resolved by our grid of runs. The width of the transition seen in the contour plot on the right panel roughly corresponds to the resolution in η of our grid of simulations and could not be any narrower given the interpolation used to generate the contours.

8. Discussion and Conclusions

8.1. Conclusions

We now review and discuss the principle conclusions of our work. We can summarize our conclusions as follows:

1. When diffusive stellar motions are accounted for, they contribute to the Galaxy becoming fully settled in a time less than, or at very least comparable to, its present age, even for slow or infrequent interstellar probes.
2. If a settlement front forms, all settleable systems behind it become filled-in in a time less than the current age of the Galaxy.
3. While settlement wave crossing and fill-in times are short, consideration of finite civilization lifetimes in a steady-state model allows for conditions in which the settled fraction X is less than 1. Thus the Galaxy may be in a steady state in which not every settleable system is currently settled.
4. Even for regions of parameter space in which one might expect $X \sim 0$ for typical regions of the Galaxy, statistical fluctuations in local density of settleable systems allows for the formation of settlement clusters which can continually resettle one another. These clusters are then surrounded by large unsettled regions. If such conditions represent the situation in our region of the Galaxy and Earth was not in one of the resettlement clusters it would be highly probable that we would not have been settled (or visited) by another civilization for some time.
5. By consideration of the convolution of steady-state solutions with geologic evidence horizons, it is possible to find situations in which Earth may not have

experienced a settlement event for longer than some horizon time (set to 1 Myr in this work) even though the Galaxy supports a population of interstellar civilizations.

Our first conclusion shows that if diffusive stellar motions are accounted for it appears almost unavoidable that if any interstellar space-faring civilization arises, the Galaxy will become fully settled in a time less than, or at very least comparable, to its present age. In particular, we confirm that thermal motions of stars prevent settlement fronts from stalling for timescales longer than the age of the Milky Way, as suggested by Brin (1983), Ashworth (2012), and Wright et al. (2014). Thus if the practical and technological impediments to interstellar settlement are overcome, then the wave of settlement should sweep across the entire Galaxy.

Note that we find that the settlement front crossing time and fill-in time takes of an order of $\lesssim 1$ Gyr even for slow probes (30 km s^{-1}). This speed is significant because it corresponds to typical interplanetary probe speeds we can design today, and is of the order of the speeds a ship of any size can achieve via gravitational slingshots with giant planets in 1 au orbits.

Our conclusions on the settlement time of the Galaxy are almost certainly lower limits for two reasons. The first is that we have not included the effects of Galactic shear or halo stars in our simulations, and these will provide additional opportunities for mixing in the case of slow ships that will cause the settlement front to expand faster than the speeds of the ships themselves.

The second is that we have assumed zero variation or improvement with time in probe launch rates, probe ranges, probe speeds, or exo-civilization lifetimes. A more realistic description of spaceflight technology on gigayear timescales would include variation among the settlements, and the expansion would likely be dominated by the high expansion rate tail of this distribution.

Our third conclusion concerns steady states for Galactic settlement. Allowing for civilizations to have finite lifetimes, the steady-state fraction of settled worlds will be a function of both civilization lifetime and the rate of encounters with empty unsettled worlds. Our results show that there are regions of the parameter space where $0 < X < 1$. Thus our steady-state model quantifies a possibility not generally considered in previous discussions of the Fermi Paradox and Galactic settlement. It is possible to achieve a galaxy in which there remain unsettled worlds even though the settlement front has crossed the entire Galactic disk.

Our fourth conclusion comes from numerical simulations of the steady-state model. Here we find under some conditions the encounter times can be so long that we would expect $X = 0$ throughout the Galaxy. This would occur because civilizations would always die out before a settlement opportunity occurred. Our simulations show however that local statistical overabundances of settleable systems can occur. This leads to clusters of closely packed settled systems surrounded by larger unsettled voids. If Earth were to exist within one of these voids then it would mean that there was a high probability that Earth might never have experienced a settlement event.

Our final conclusion concerns the temporal aspect of the Fermi Paradox and Hart’s Fact A—the lack of any obvious settlement of the solar system—which Hart argues compels the conclusion that there can be no other technological civilizations in the Galaxy. By including a finite time horizon past which evidence of prior settlement civilization might not be seen, we

have shown that it is possible to break the link between Hart’s Fact A and his conclusion. To wit, it is possible to have a Galaxy with some non-zero settlement fraction and still have evidence for a prior Earth visit lie over the horizon available via Earth’s geological record. Indeed, the last three of our conclusions all break the link between conclusions about rapid Galactic settlement and the current absence of technological civilizations. This occurs because the steady-state model implies that not all settleable systems need be currently occupied.

8.2. Discussion

For low densities of settleable systems, our steady-state calculation finds that consistency with the lack of evidence for Earth’s past settlement requires that each civilization has, on average, only one chance to reproduce, i.e., to settle another world. This does not mean that only one settlement probe was launched however. Our result can be interpreted to mean that on average only one settlement probe was successful. Inherent in our calculation was the settlement fraction f . Our steady-state calculation was carried out in the low density limit which implies $f < 1$ (note that in the high density limit the fraction simply tends to $X = 1$). As described earlier, low values of f implies that good planets are hard find. Our steady-state calculations in the low density limit further imply that successful settlements are hard to achieve. The lack of settlement success could come for many reasons ranging from failure of interstellar vessels capable of establishing persistent settlements to the inability to develop viable progeny civilizations on new worlds.

Because Hart’s conclusions stem from his assumption that we would have noticed if extraterrestrial technology had ever settled the solar system, they are challenged by the work of Freitas (1985), Schmidt & Frank (2019), Davies (2012), and Haqq-Misra & Kopparapu (2012) who show that this is not necessarily the case.

We can go further though: Hart’s conclusions are also subject to the assumption that the solar system would be considered settleable by any of the exo-civilizations it has come within range of. The most extravagant contradiction of this assumption is the Zoo Hypothesis (Ball 1973), but we need not invoke such solipsist positions (Sagan & Newman 1983) to point out the flaw in Hart’s reasoning here. One can imagine many reasons why the solar system might not be settleable (i.e., not part of the fraction f in our analysis), including the *Aurora* effect mentioned in Section 1 or the possibility that they avoid settling the environment near the Earth exactly because it is inhabited with life.

In particular, the assumption that the Earth’s life-sustaining resources make it a particularly *good* target for extraterrestrial settlement projects could be a naive projection onto exo-civilizations of a particular set of human attitudes that conflate expansion and exploration with conquest of (or at least indifference toward) native populations (Wright & Oman-Reagan 2018). One might just as plausibly posit that any extremely long-lived civilization would appreciate the importance of leaving native life and its near-space environment undisturbed.

So our conclusions have strong implications for the likelihood of success of SETI, but the specific nature of that optimism is strongly dependent on assumptions regarding either the limits of technology or the agency of exo-

civilizations. If large-scale terraforming is not a realistic possibility then f may be limited by worlds matching the biology of the parent or Ur civilization. Earth may simply not be one of those worlds even though such settlements exist elsewhere. Likewise if, for whatever reason, all extraterrestrial civilizations that have had to the opportunity to settle the solar system have avoided it, then the Milky Way should be filled with stars hosting potentially detectable extraterrestrial technology if even a single settlement front has never been established.

If instead one follows Hart and assumes that solar system settlements would be inevitable, then our analysis quantifies the regions of parameter space consistent with his Fact A in which the Galaxy is filled with, or devoid of, space-faring technology. Those sets of parameters in which Fact A does imply an empty Milky Way are those in which Galactic settlement is especially efficient. This implies that *other* galaxies where technological life has arisen should have been thoroughly settled, raising the prospects that they might be detected at extragalactic distances.

We note also that in the classic argument, Hart’s Fact A is linked to conclusions about exo-civilizations because it is assumed that interstellar travel is a natural result of their evolution. But this need not be the case. In Ashworth (2012), the energetics of developing interstellar ships that could host long-term viable populations was explored. Given the travel times between stars, these would be multigenerational world ships. Ashworth (2012) attempted to calculate the cost of building such machines, including factors such as speed and mass. His finding was that that economies equivalent to that of entire solar systems would be required to develop and launch world ships. As an example, consider his “medium multi-generational cruiser” case. This was ship traveling at $v = 0.05c$, carrying a population of 10^4 people and weighing 10^7 tonnes. Such a ship would require a power of 6900 zettajoules (ZJ). He estimates that a solar-system-wide civilization of 900 billion people would generate 1136 ZJ per year. Thus while the creation of world ship by such an economy would be possible, it would require a significant proportion the civilization’s resources. We note that these estimates are, of course, highly speculative and Ashworth (2012) also provides estimates for solar-system-wide civilizations generating even higher power economies.

For our present results these factors indicate that it is possible that developing the requirements for interstellar settlement may be expensive enough to be universally prohibitive. In addition, if establishment of viable settlements proves difficult, meaning the success rate of world ships is low, then civilizations may be unwilling to continue investing in them over time. This is particularly true if one considers that the long travel and communication times may make it difficult to establish an interstellar civilization. Unless the individuals in the species driving the settlement have very long lifetimes (>100 yr) it is difficult to see how a galactic scale culture can arise (i.e., commerce, etc.; Krugman 2010). Thus each settlement may, in practice, be relatively isolated culturally, which may limit the effort civilizations are willing to put into long-term programs of expansion.

It is also worth considering the distribution of natural catastrophes which might lead to end of settlements. Weak constraints might be obtained using Earth as an example by looking for cross-correlation with the ages of the impact craters, super-volcanic deposits, and extinctions. The

timescales between such events is likely longer than 1 Myr and more work can be done to explore the question of look-back horizons for evidence of settlement events. As noted earlier, as the look-back time is extended plots like those shown in Figure 11 tend to have smaller intermediate probabilities for short-lived settlements.

In summary: our work demonstrates that even though settlement fronts can be expected to cross the Galaxy quickly, every settleable system need not be inhabited. We note that much work needs to be done to extract the maximum amount of information from this fact when convolved with both the expected conditions for different regions of the Galaxy along with what can plausibly be expected from Earth’s geologic record. In particular further studies of the settlement steady states may help understand the creation, extent, and longevity of settlement voids, which provide one explanation for the lack of evidence for Earth’s past settlement. Our calculations open a new avenue in consideration of exo-civilizations and their prevalence in the Galaxy, and have strong but assumption-dependent implications for the prospects of the success of SETI.

We thank Milan M. Ćirković Woody Sullivan, Gregory Van Maanen, Andrew Siemion, and David Brin for the helpful discussions. We also acknowledge the New York Academy of Sciences for hosting a panel discussion which helped create the impetus for this work. C.S. acknowledges support from the NASA Astrobiology Program through participation in the Nexus for Exoplanet System Science and NASA Grant NNX15AK95G. The Center for Exoplanets and Habitable Worlds is supported by the Pennsylvania State University, the Eberly College of Science, and the Pennsylvania Space Grant Consortium. The authors thank the Center for Integrated Research Computing (CIRC) at the University of Rochester for providing computational resources and technical support.

Appendix

A.1. Calculating $\mathcal{V}(n)$

The expectation value for the maximum velocity from N samples from a Maxwell–Boltzmann distribution

$$\mathcal{V}(n) = E[V_{\max}(n)], \quad (32)$$

where V_{\max} is the distribution of maximum values from a sample of n velocities taken from a velocity distribution $P(v)$. For a given velocity distribution $P(v)$ with CDF $C(v)$, the odds that n random samples are below a particular value v_0 is $C(v_0)^n$. This is then the CDF for the maximum value is $C_{\max}(v)$. To get the expectation value, we have $E[v] = \int v \times P_{\max}(v) dv$, where

$$P_{\max} = \frac{dC_{\max}(v)}{dv} = \frac{dC(v)^n}{dv}. \text{ So we have}$$

$$\mathcal{V}(n) = \int v \frac{d(C(v)^n)}{dv} dv. \quad (33)$$

A.2. Including Targeted Systems

If we assume that settled systems do not launch probes at other systems that are already settled, we can modify our model by including the fraction of systems that are targeted in addition

to those that are settled and habitable,

$$\frac{dN_s}{dt} = \frac{1}{t_p}N_t - \frac{1}{T_s}N_s \quad (34)$$

$$\frac{dN_t}{dt} = \frac{1}{T_l}N_sN_h - \frac{1}{t_p}N_t \quad (35)$$

$$\frac{dN_h}{dt} = \frac{1}{T_s}N_s - \frac{1}{T_l}N_sN_h, \quad (36)$$

where N_s is the number of settled systems, N_t is the number of targeted systems, and N_h is the number of habitable systems. In addition t_p is the probe travel time, T_s is the settled time, and T_l is the effective launch time (due to either probe production rates or encounter times). We can dimensionalize the equation where $X = \frac{N_s}{N}$, $Y = \frac{N_t}{N}$ and $\frac{N_h}{N} = 1 - X - Y$. The resulting equations are

$$\frac{dX}{dt} = \frac{Y}{t_p} - \frac{1}{T_s}X \quad (37)$$

$$\frac{dY}{dt} = \frac{1}{T_l}X(1 - X - Y) - \frac{Y}{t_p}. \quad (38)$$

Solving for equilibrium, we find $X = \frac{T_s - T_l}{T_s + t_p}$

ORCID iDs

Jonathan Carroll-Nellenback  <https://orcid.org/0000-0003-3265-7210>

Jason Wright  <https://orcid.org/0000-0001-6160-5888>

References

- Ashworth, S. 2012, *JBIS*, **65**, 140
 Ashworth, S. 2014, *JBIS*, **67**, 224
 Ball, J. A. 1973, *Icar*, **19**, 347
 Bjørk, R. 2007, *IJAsB*, **6**, 89
 Bracewell, R. N. 1960, *Natur*, **186**, 670
 Brin, G. D. 1983, *QJRAS*, **24**, 283
 Chyba, C. F., & Hand, K. P. 2005, *ARA&A*, **43**, 31
 Cirkovic, M. M. 2018, *The Great Silence: Science and Philosophy of Fermi's Paradox* (Oxford: Oxford Univ. Press)
 Cotta, C., & Morales, A. 2009, *JBIS*, **62**, 82
 Davies, P. C. W. 2011, *The Eerie Silence: Renewing Our Search for Alien Intelligence* (Orlando, FL: Mariner Books)
 Davies, P. C. W. 2012, *AcAau*, **73**, 250
 Drake, F. D. 1965, in *Current Aspects of Exobiology*, ed. G. Mamikunian & M. H. Briggs (Oxford: Oxford Univ. Press), 323
 Eric, M. 1985, *Where Is Everybody?: An Account of Fermi's Question* (Los Alamos, NM: Los Alamos National Laboratory)
 Forgan, D. H. 2009, *IJAsB*, **8**, 121
 Frank, A. 2018, *Light of the Stars. Alien Worlds and The Fate of the Earth* (New York: W.W. Norton)
 Freitas, R. A. 1980, *JBIS*, **33**, 251
 Freitas, R. A. 1985, *Icar*, **62**, 518
 Galera, E., Galanti, G., & Kinouchi, O. 2018, *IJAsB*, **18**, 316
 Gray, R. 2015, *AsBio*, **15**, 3195
 Gros, C. 2005, *JBIS*, **58**, 108
 Hair, T. W., & Hedman, A. D. 2013, *IJAsB*, **12**, 45
 Haqq-Misra, J., & Kopparapu, R. K. 2012, *AcAau*, **72**, 15
 Haqq-Misra, J. D., & Baum, S. D. 2009, *JBIS*, **62**, 47
 Hart, M. H. 1975, *QJRAS*, **16**, 128
 Henry, T. J., Jao, W.-C., Winters, J. G., et al. 2018, *AJ*, **155**, 265
 Jones, E. M. 1976, *Icar*, **28**, 421
 Jones, E. M. 1978, *JBIS*, **31**, 103
 Jones, E. M. 1981, *Icar*, **46**, 328
 Kinouchi, O. 2001, arXiv:condmat/0112137
 Krugman, P. 2010, *Economic Inquiry*, **48**, 1119
 Kuiper, T. B. H., & Morris, M. 1977, *Sci*, **196**, 616
 Landis, G. A. 1998, *JBIS*, **51**, 163
 Lin, H. W., & Loeb, A. 2015, *ApJL*, **810**, L3
 Newman, W. I., & Sagan, C. 1981, *Icar*, **46**, 293
 Robinson, K. S. 2015, *Aurora* (London: Orbit Books)
 Sagan, C., & Newman, W. I. 1983, *QJRAS*, **24**, 113
 Schmidt, G., & Frank, A. 2019, *IJAsB*, **18**, 142
 Siemion, A. P. V., Demorest, P., Korpela, E., et al. 2013, *ApJ*, **767**, 94
 Stull, M. A. 1979, *JBIS*, **32**, 221
 Tarter, J. 2001, *ARA&A*, **39**, 511
 Tarter, J. C. 2007, *HiA*, **14**, 14
 Tarter, J. C., Agrawal, A., Ackermann, R., et al. 2010, *Proc. SPIE*, **7819**, 781902
 Tipler, F. J. 1980, *QJRAS*, **21**, 267
 von Hoerner, S. 1975, *JBIS*, **28**, 691
 Vukotić, B., & Ćirković, M. M. 2012, *OLEB*, **42**, 347
 Walters, C., Hoover, R. A., & Kotra, R. K. 1980, *Icar*, **41**, 193
 Webb, S. 2002, *If the Universe is Teeming with Aliens...Where is Everybody?: Fifty Solutions to the Fermi Paradox and the Problem of Extraterrestrial Life* (New York: Praxis Book/Copernicus Books)
 Wiley, K. B. 2011, arXiv:1111.6131
 Wright, J. T. 2018, *IJAsB*, **17**, 96
 Wright, J. T., Kanodia, S., & Lubar, E. 2018, *AJ*, **156**, 260
 Wright, J. T., Mullan, B., Sigurdsson, S., & Povich, M. S. 2014, *ApJ*, **792**, 26
 Wright, J. T., & Oman-Reagan, M. P. 2018, *IJAsB*, **17**, 177
 Zackrisson, E., Calissendorff, P., Asadi, S., & Nyholm, A. 2015, *ApJ*, **810**, 23