Predictions of the *WFIRST* Microlensing Survey. I. Bound Planet Detection Rates

Matthew T. Penny¹ , B. Scott Gaudi¹ , Eamonn Kerins² , Nicholas J. Rattenbury³ , Shude Mao^{2,4,5} , Annie C. Robin⁶, and Sebastiano Calchi Novati⁷

¹ Department of Astronomy, The Ohio State University, 140 West 18th Avenue, Columbus, OH 43210, USA

² Jodrell Bank Centre for Astrophysics, Alan Turing Building, University of Manchester, Manchester M13 9PL, UK

³ Department of Physics, University of Auckland, Private Bag 92019, Auckland, New Zealand

⁴ Physics Department and Tsinghua Centre for Astrophysics, Tsinghua University, Beijing 100084, People's Republic of China

⁵ National Astronomical Observatories, Chinese Academy of Sciences, A20 Datun Road, Chaoyang District, Beijing 100012, People's Republic of China

⁶ Institut Utinam, CNRS UMR 6213, OSU THETA, Université Bourgogne-Franche-Comté, 41bis avenue de l'Observatoire, F-25000 Besançon, France

⁷ IPAC, Mail Code 100-22, Caltech, 1200 East California Boulevard, Pasadena, CA 91125, USA

Received 2018 August 6; revised 2018 December 21; accepted 2018 December 31; published 2019 February 25

Abstract

The *Wide Field InfraRed Survey Telescope* (*WFIRST*) is the next NASA astrophysics flagship mission, to follow the *James Webb Space Telescope*. The *WFIRST* mission was chosen as the top-priority large space mission of the 2010 astronomy and astrophysics decadal survey in order to achieve three primary goals: to study dark energy via a wide-field imaging survey, to study exoplanets via a microlensing survey, and to enable a guest observer program. Here we assess the ability of the several *WFIRST* designs to achieve the goal of the microlensing survey to discover a large sample of cold, low-mass exoplanets with semimajor axes beyond roughly one astronomical unit, which are largely impossible to detect with any other technique. We present the results of a suite of simulations that span the full range of the proposed *WFIRST* architectures, from the original design envisioned by the decadal survey, to the current design, which utilizes a 2.4 m telescope donated to NASA. By studying such a broad range of architectures, we are able to determine the impact of design trades on the expected yields of detected exoplanets. In estimating the yields we take particular care to ensure that our assumed Galactic model predicts microlensing event rates that match observations, consider the impact that inaccuracies in the Galactic model might have on the yields, and ensure that numerical errors in light-curve computations do not bias the yields for the smallest-mass exoplanets. For the nominal baseline *WFIRST* design and a fiducial planet mass function, we predict that a total of ~ 1400 bound exoplanets with mass greater than $\sim 0.1 M_{\oplus}$ should be detected, including ~ 200 with mass $\lesssim 3 M_{\oplus}$. *WFIRST* should have sensitivity to planets with mass down to $\sim 0.02 M_{\oplus}$, or roughly the mass of Ganymede.

Key words: gravitational lensing: micro – planets and satellites: detection – space vehicles: instruments

1. Introduction

The study of the demographics of exoplanets, the end result of the planet formation process, has entered a statistical age. Large samples of transiting planets from *Kepler* (e.g., Thompson et al. 2018), massive planets at small to moderate separations from ground-based radial velocity (RV) surveys of planetary systems in the solar neighborhood (e.g., Udry & Santos 2007; Winn & Fabrycky 2015), and direct imaging studies of young planets at large separations (e.g., Bowler 2016), are beginning to reveal the complex distribution of exoplanets as a function of mass and separation from their host stars, and the properties of the host stars themselves.

Data from the *Kepler* mission have revealed a sharp rise in the occurrence rate of hot and warm planets as radius decreases down to about $2.8 R_{\oplus}$, before leveling off (e.g., Howard et al. 2010; Fressin et al. 2013). Precise spectroscopic measurements of *Kepler*'s super-Earth hosts have revealed a radius dichotomy between large and small super-Earths (Fulton et al. 2017), which is likely due to atmospheric stripping (Owen & Wu 2017). At large orbital distances $\gtrsim 10$ au, direct imaging searches have found young, massive planets to be present, but rare (e.g., Nielsen & Close 2010; Bowler et al. 2015; Chauvin et al. 2015). However, there remains a large area of the exoplanet parameter space—orbits beyond ~ 1 au and masses less than that of Jupiter—that remains relatively unexplored by transit, RV, and direct imaging techniques.

Indeed, if every planetary system resembled our own, only a handful of planets would have been discovered by the RV, transit, or direct imaging techniques to date. This fact begs the question: is our solar system architecture rare? If so, why?

To obtain a large sample of exoplanets beyond 1 au and across a large range of masses requires a different technique. Gravitational microlensing enables a statistical survey of exoplanet populations beyond 1 au, because its sensitivity peaks at the Einstein radius of its host stars (Mao & Paczynski 1991; Gould & Loeb 1992; Bennett & Rhie 1996). For stars along the line of sight to the Galactic bulge (where the microlensing event rate is highest) the physical Einstein radius is typically 2–3 au (see, e.g., the review of Gaudi 2012). Thanks to the fact that microlensing is sensitive directly to a planet's mass and not its light or effect on a luminous body, the technique's sensitivity extends out to all orbital radii beyond ~ 1 au (Bennett & Rhie 2002).

Perhaps the most important reason to perform a large exoplanetary microlensing survey is that it opens up a large new region of parameter space. The history of exoplanet searches has been one of unexpected discoveries. At every turn, when a new area of parameter space has been explored, previously unexpected planetary systems have been found. This process began with the pulsar planets (Wolszczan & Frail 1992) and relatively short-period giant planets discovered by the first precision RV searches sensitive to planets (Campbell et al. 1988; Latham et al. 1989; Mayor & Queloz 1995). As RV surveys' sensitivities and durations grew, highly eccentric

massive planets and low-mass Neptunes and super-Earth planets were discovered (e.g., Naef et al. 2001; Butler et al. 2004; Rivera et al. 2005). When originally conceived, and with the solar system as a guide, the *Kepler* mission aimed to detect potentially habitable Earth-sized planets in ~ 1 au orbits around solar-like stars (Borucki et al. 2003). Were all exoplanet systems like our own, *Kepler* would have found few or no planets (Burke et al. 2015), and those that it did would have been at the limit of its signal-to-noise ratio (S/N). This result was obviously pre-empted by the discovery of hot Jupiters, which demonstrated conclusively that *not* all planetary systems have architectures like our own. *Kepler* itself has gone on to discover thousands of planetary systems very unlike ours, including tightly packed multiplanet systems (e.g., Lissauer et al. 2011) and circumbinary planets (e.g., Doyle et al. 2011), to name but a few examples. Even moving into unprobed areas of the host mass parameter space has revealed unexpected systems such as TRAPPIST-1 (Gillon et al. 2016) and KELT-9 (Gaudi et al. 2017). Direct imaging searches have revealed young, very massive planets that orbit far from their hosts (e.g., Chauvin et al. 2004), the most unusual (from our solar system-centric viewpoint) being the four-planet system around HR 8799 (Marois et al. 2008, 2010).

Despite finding the planetary system that is arguably most similar to our own (the OGLE-2006-BLG-109 Jupiter–Saturn analog; Gaudi et al. 2008; Bennett et al. 2010b), ground-based microlensing surveys too have discovered unexpected systems. For example, microlensing searches have found several massive planets around M-dwarf stars (e.g., Dong et al. 2009) that appear to at least qualitatively contradict the prediction of the core accretion theory that giant planets should be rare around low-mass stars (Laughlin et al. 2004). Measurements of planet occurrence rates from microlensing also superficially appear to contradict previous RV results, although a more careful analysis indicates that the microlensing and RV results are consistent (Clanton & Gaudi 2014a, 2014b; Montet et al. 2014). Other notable microlensing discoveries include circumbinary planets (Bennett et al. 2016), planets on orbits of ~ 1 –10 au around components of moderately wide binary stars (e.g., Gould et al. 2014b; Poleski et al. 2014), planets on wide orbits comparable to Uranus and Neptune (e.g., Poleski et al. 2017), and planets orbiting ultracool dwarfs (e.g., Shvartzvald et al. 2017).

Having spent over a decade conducting two-stage survey-plus-follow-up planet searches (see, e.g., Gould et al. 2010, for a review), microlensing surveys have entered a second-generation mode, which relies only on survey observations. The Optical Gravitational Lensing Experiment (OGLE; Udalski et al. 2015a), Microlensing Observations in Astrophysics (MOA; Sako et al. 2007), and three Korea Microlensing Telescope Network (KMTNet) telescopes (Kim et al. 2016) span the southern hemisphere and provide continuous high-cadence microlensing observations over tens of square degrees every night that weather allows. Such global, second-generation, pure survey-mode microlensing surveys will enable the initial promise of microlensing to provide the large statistical samples of exoplanets necessary to study demographics (Henderson et al. 2014a), and have begun to deliver (Shvartzvald et al. 2016; Suzuki et al. 2016). It has long been recognized (e.g., Bennett & Rhie 2002; Beaulieu et al. 2008; Gould 2009) that exoplanet microlensing surveys are best conducted from space, thanks to the greater

ability to resolve stars in crowded fields and to continuously monitor fields without interruptions from weather or the day–night cycle.

The *Wide Field InfraRed Survey Telescope* (WFIRST) (Spergel et al. 2015) is a mission conceived by the 2010 decadal survey panel (Committee for a Decadal Survey of Astronomy & Astrophysics 2010) as its top-priority large astrophysics mission. It combines mission proposals to study dark energy with weak lensing, baryon acoustic oscillations, and supernovae (Joint Dark Energy Mission-Omega, JDEM-Omega; Gehrels 2010), with a gravitational microlensing survey (Microlensing Planet Finder, MPF; Bennett et al. 2010a), a near-infrared sky survey (Near Infrared Sky Surveyor; Stern et al. 2010) and a significant guest observer component (Committee for a Decadal Survey of Astronomy & Astrophysics 2010). The later addition of a high-contrast coronagraphic imaging and spectroscopic technology demonstration instrument (Spergel et al. 2015) addresses a top medium-scale 2010 decadal survey priority as well. In this paper we examine only the microlensing survey component of the mission.

The structure of the paper is as follows. In Section 2 we describe the WFIRST mission, and each of its design stages. In Section 3 we describe the simulations we have performed. Section 4 presents the yields of the baseline simulations, while Section 5 considers the effects of various possible changes to the mission design. Section 6 discusses the uncertainties that affect our results and how they might be mitigated by future observations, modeling, and simulations. Section 7 gives our conclusions.

2. WFIRST

2.1. Goals of the WFIRST Microlensing Survey

A primary science objective of the WFIRST mission is to conduct a statistical census of exoplanetary systems, from 1 au out to free-floating planets, including analogs to all of the solar system planets with masses greater than Mars, via a microlensing survey. It is in the region of ~ 1 –10 au that the microlensing technique is most sensitive to planets over a wide range of masses (Mao & Paczynski 1991; Gould & Loeb 1992), and where other planet detection techniques lack the sensitivity to detect low-mass planets within reasonable survey durations or present-day technological limits. However, the 1–10 au region is perhaps the most important region of protoplanetary disks and planetary systems for determining their formation and subsequent evolution, and can have important effects on the habitability of planets.

The enhancement in surface density of solids at the water ice line, ~ 1.5 –4 au from the star, is thought to be critical for the formation of giant planets (Hayashi 1981; Ida & Lin 2004; Kennedy et al. 2006). Nevertheless, not all stars produce giant planets that survive (e.g., Winn & Fabrycky 2015). It remains to be seen whether this is due to inefficient production of giant planets, or a formation process that is $\sim 100\%$ efficient followed by an effective destruction mechanism, such as efficient disk migration (e.g., Goldreich & Tremaine 1980), or ejection or host star collisions caused by dynamics (e.g., Rasio & Ford 1996). In the core accretion scenario (e.g., Goldreich & Ward 1973; Pollack et al. 1996), runaway gas accretion onto protoplanet cores to produce giant planets is an inevitability, if the core manages to accrete a gaseous envelope as massive as

the core before the gas dissipates from the protoplanetary disk (Mizuno 1980). If core growth rate is the rate-limiting step in the production of giant planets, and the process is indeed inefficient, then we can expect a population of “failed cores” of various masses with a distribution that peaks near the location of the ice line. Conversely, if giant planet formation and subsequent destruction is efficient, we can expect the formed giant planets to clear their orbits of other bodies, and thus would expect to see a deficit of low-mass planets in *WFIRST*’s region of sensitivity.

It is clear that planetary systems are not static, and the orbits of planets can evolve during and after the planet formation process, first via drag forces while the protoplanetary disk is in place (Lin et al. 1996), and subsequently by N -body dynamical processes once the damping effect of the disk is removed (Rasio & Ford 1996). In addition to rearranging the orbits of planets that remain bound, the chaotic dynamics of multiplanet systems can result in planets being ejected (e.g., Safronov 1972). The masses and number of ejected planets from the system will be determined by the number of planets in their original systems, their masses, and orbital distribution (e.g., Papaloizou & Terquem 2001; Chatterjee et al. 2008; Jurić & Tremaine 2008; Barclay et al. 2017). Thus the mass function of ejected, or free-floating, planets can be an important constraint on the statistics of planetary systems as a whole. Because they emit very little light, only microlensing observations can be used to detect rocky free-floating planets. For masses significantly below Earth’s, only space-based observations can provide the necessary combination of photometric precision, cadence, and total number of sources monitored in order to collect a significant sample of events.

For both bound and free-floating objects it is valuable to extend the mass sensitivity of an exoplanet survey down past the characteristic mass scales of planet formation theory, where the growth behavior of forming planets changes. This is particularly the case for boundaries in mass between low and high growth rates, as these should be the locations of either pile-ups or deficits in the mass function, depending on the sense of the transition. In the core accretion scenario of giant planet formation, moving from high to low masses, characteristic mass scales include the critical core mass for runaway gas accretion at $\sim 10 M_{\oplus}$ (Mizuno 1980), the isolation mass of planetary embryos at $\sim 0.1 M_{\oplus}$ (Kokubo & Ida 2002), and the transition core mass for pebble accretion at $\sim 0.01 M_{\oplus}$ (Lambrechts & Johansen 2012). The detection of features due to these characteristic mass scales would be strong evidence in support of current planet formation theory. Additionally, a statistical accounting of planets on wider orbits more generally will be a valuable test of models of planet formation developed to explain the large occurrence rate of super-Earths closer than 1 au.

An estimate of the occurrence rate of rocky planets in the habitable zones (e.g., Kasting et al. 1993; Kopparapu et al. 2013) of solar-like stars, $\eta_{\oplus}$, is an important ingredient for understanding the origins and evolution of life, and the uniqueness of its development on Earth. However, it is precisely this location where it is both hardest to detect $\sim 1 M_{\oplus}$ planets around $\sim 1 M_{\odot}$ stars, while also remaining tantalizingly achievable. Tiny signals recurring on approximately year timescales mean that transit, RV, and astrometric searches must run for multiple years to make robust detections. Only the transit technique has demonstrated the necessary

precision to date and, even so, *Kepler* fell just short of the mission duration necessary to robustly measure $\eta_{\oplus}$ (Burke et al. 2015). Direct imaging of habitable exoplanets will require significant technological advances (e.g., Mennesson et al. 2016; Bolcar 2017; Wang et al. 2018), several of which *WFIRST*’s coronagraphic instrument will demonstrate, in order to reach the contrast and inner working angles required, and the observing time required to perform a blind statistical survey to measure $\eta_{\oplus}$ may be prohibitively expensive. The typical inner sensitivity limit for a space-based microlensing survey, which is proportional to the host mass $M^{1/2}$, crosses the habitable zone, which scales as $\sim M^{3.5}$, at $\sim 1 M_{\odot}$. Nevertheless the detection efficiency for low-mass-ratio planets inside the Einstein ring ($R_E \sim 3$ au) is very small, and solar-mass stars make up only a small fraction of the lens population, so large, long-duration microlensing surveys from space are required to robustly measure $\eta_{\oplus}$ with microlensing. In all likelihood, no single technique will prove sufficient, and it will be necessary to combine measurements from multiple techniques to be confident in the accuracy of the $\eta_{\oplus}$ determination. If the habitable zone is extended outward (e.g., Seager 2013), by volcanic outgassing of H_2 (Ramirez & Kaltenegger 2017) or some other process, the number of habitable planets that space-based microlensing searches are sensitive to increases significantly.

Each of the goals described above can be addressed in whole or in part by studying the statistics of a large sample of planets with orbits in the range of 1–10 au, and a similar sample of unbound planets. Such a sample can only be delivered by a space-based microlensing survey. Astrometry from *Gaia* can be used to discover a large sample of giant planets in similar orbits, but it will not have the sensitivity to probe below $\sim 30 M_{\oplus}$ (Perryman et al. 2014). Space-based transit surveys have sensitivity to very small and thus low-mass planets, but would be required to observe for decades to cover the same range of orbital separations as does microlensing. Ground-based microlensing searches have sensitivity to low-mass exoplanets down to $\sim 1 M_{\oplus}$, as recently demonstrated by Bond et al. (2017) and Shvartzvald et al. (2017), but are limited from gathering large samples of such planets or extending their sensitivity to masses significantly smaller than this by a combination of the more limited photometric precision possible from the ground, the larger angular diameter of source stars for which high precision is possible, and the lower density of such sources on the sky.

A critical element in measuring the mass function of planets from microlensing events is actually measuring planets’ masses. The light curve of a binary microlensing event alone only reveals information about the mass *ratio* q , unless the event is long enough to measure the effect of annual microlensing parallax on the light curve, or the event is observed from two widely separated observers, e.g., a spacecraft such as *Spitzer* (e.g., Udalski et al. 2015b) or *Kepler* (see Henderson et al. 2016 for a review). It is also necessary to measure finite-source effects in the light curve, but this is routinely achieved for almost all planetary microlensing events observed to date, and will likely be possible for the majority of *WFIRST*’s planetary events (e.g., Zhu et al. 2014). High-resolution imaging enables an alternative method to measure the host and planet mass. Over time, the source and lens star involved in a microlensing event will move apart, and the lens, if bright enough, will be detectable, either as an elongation in the combined source–lens image, as a shift in the centroid of

the pair as a function of color, or as a resolved star if its proper motion relative to the source is large enough (e.g., Bennett et al. 2007). The measured separation between the stars, and the color and magnitude of the lens star, can be combined with the measurement of the event timescale to uniquely determine the mass of the lens. A principal requirement of the *WFIRST* mission is the ability to make these measurements routinely for most events. This is made possible by the resolution achievable from space, which reduces the time that one has to wait for a source and lens to separate and provides access to fainter sources that do not outshine faint, main-sequence lens stars. For *WFIRST* to make these measurements it is necessary that it observe the microlensing fields over a time baseline of four or more years.

In this paper we will only address *WFIRST*'s ability to measure the mass function of bound planets. The challenges and opportunities to detect free-floating planets, and planets in the habitable zone, differ somewhat from those for the general bound planet population. We have therefore elected to give them the full attention that they deserve in subsequent papers, rather than provide only the limited picture that would be possible in this paper.

2.2. Evolution of the WFIRST Mission: Design Reference Missions, AFTA, and Cycle 7

The *WFIRST* mission is in the process of ongoing design refinement, and has gone through four major phases so far. This paper presents analysis of each of these missions, even though some of these designs are no longer under active development. This is important for two reasons. First, we are documenting the quantitative simulations that have informed the *WFIRST* microlensing survey design process from the first science definition team (SDT). Second, each design represents a internally self-consistent set of mission design parameters which, when evolved to a new mission design, necessarily captures the majority of covariance between all of the possible design choices. These covariances are difficult to account for in simulations that might aim to investigate variations in individual parameters by isolating their effects.

The first *WFIRST* design, the Interim Design Reference Mission (IDRM) was based directly on the *WFIRST* mission proposed by the decadal survey and described in the first *WFIRST* SDT's interim report (Green et al. 2011). This in turn was based on the design for the JDEM-Omega mission (Gehrels 2010). The IDRM consisted of an unobstructed 1.3 m telescope with a 0.294 deg^2 near-infrared imaging channel with broadband filters spanning $\sim 0.76\text{--}2.0 \mu\text{m}$, including a wide $1\text{--}2 \mu\text{m}$ filter for the microlensing survey, and two slitless spectroscopic channels.

The final report of the first *WFIRST* SDT (Green et al. 2012) presented two Design Reference Missions (DRM1 and DRM2). DRM1 was an evolution of the IDRM, adhering to the recommendation of the decadal survey to only use fully developed technologies. It improved on the IDRM by increasing the upper wavelength cutoff of the detectors to $2.4 \mu\text{m}$, and removing the two spectroscopic channels. The detectors and prism elements were added to the imaging channel to increase its field of view to 0.377 deg^2 .

DRM2 was a design intended to reduce the cost of the mission. This was done by reducing the size of the primary mirror to 1.1 m (in order to fit onto a less costly launch vehicle). It also switched to a larger format $4\text{k} \times 4\text{k}$ detector to reduce

the number of detectors while increasing the field of view to 0.585 deg^2 , at the cost of additional detector development. The larger field of view also allowed the mission duration to be shortened to three years instead of five.

The *WFIRST* design process was disrupted in 2012 when NASA was gifted two 2.4 m telescope mirrors and optical tube assemblies by another government agency. The value of these telescopes to the *WFIRST* mission was initially assessed in a report by Dressler et al. (2012), and the mission designed around one of the 2.4 m telescopes was dubbed *WFIRST*-AFTA (AFTA standing for Astrophysically Focussed Telescope Assets). A new SDT was assembled to produce an AFTA DRM, which added a coronagraphic instrument channel to the mission (Spergel et al. 2013). The design of the wide field instrument (WFI) also changed, requiring a finer pixel scale to sample the smaller point-spread function (PSF) of the 2.4 m telescope, and hence a smaller field of view of 0.282 deg^2 . Unlike the previous designs, the telescope has an obstructed pupil, so the PSF has significant diffraction spikes that the previous versions did not. The design also required a shorter wavelength cutoff of $2.0 \mu\text{m}$ due to concerns about the ability to operate the telescope at the low temperature required for $2.4 \mu\text{m}$ observations. The final results of this design process were presented by Spergel et al. (2015).

WFIRST entered the formulation phase (phase A) in early 2016. The AFTA design was adopted, and the mission reverted to its simpler naming of *WFIRST*. Extensive design and testing work has been conducted since formulation began. This includes a large amount of detector development, validation of the ability of the telescope to operate at lower temperature, redesign of the WFI, and consideration of various mission scopes. Very recently, *WFIRST* entered the second phase of mission development (phase B).

The most significant update to the design affecting the microlensing survey is a more pessimistic accounting of the observatory's slew time performance compared to the AFTA design. Another important change is a rotation of the elongated detector layout by 90° . While the mission design continues to evolve, we present simulations here that most closely match the Cycle 7 design, and so throughout the paper we will refer to the design as *WFIRST* Cycle 7.

Table 1 summarizes the parameters of each mission design that we study. We use these parameters for the results presented in Sections 4 and 4.1. In Section 5 we present the results of "trade-off" simulations that were conducted during the design process, when many of the mission parameters changed regularly. Between any given set of trade-off simulations, the exact values of many of the simulation parameters changed by small amounts. Rather than tediously detail each of these parameter changes, which have little effect on the absolute yields, we will only indicate changes to parameters where they are important to each study. Invariably, these will be the independent variables of each study, or parameters closely related to them. As the unimportant parameters do not change internal to each trade study, we compute differential yield measurements, i.e., yields relative to a fiducial design.

2.3. The WFIRST Microlensing Survey

The full operations concept for the *WFIRST* mission must ensure that the spacecraft can conduct all the observations necessary to meet its primary mission requirements, while maintaining sufficient flexibility to conduct a significant

Table 1
Adopted Parameters of Each Mission Design

	IDRM		DRM1		DRM2		AFTA		WFIRST Cycle 7	
Reference	Green et al. (2011)		Green et al. (2012)		Green et al. (2012)		Spergel et al. (2015)		... ^{a,b}	
Mirror diameter (m)	1.3		1.3		1.1		2.36		2.36	
Obscured fraction (area, %)	0		0		0		13.9		13.9	
Detectors	7 × 4 H2RG-10		9 × 4 H2RG-10		7 × 2 H4RG-10		6 × 3 H4RG-10		6 × 3 H4RG-10	
Plate scale (″/pix)	0.18		0.18		0.18		0.11		0.11	
Field of view (deg ²)	0.294		0.377		0.587		0.282		0.282	
Fields	7		7		6		10		7	
Survey area (deg ²)	2.06		2.64		3.52		2.82		1.97	
Avg. slew and settle time (s)	38		38		38		38		83.1	
Orbit	L2		L2		L2		Geosynchronous		L2	
Total survey length (day)	432		432		266		411 ^c		432	
Season length (day)	72		72		72		72		72	
Seasons	6		6		3.7		6		6	
Baseline mission duration (yr)	5		5		3		6		5	
Primary bandpass (μm)	1.0–2.0 (W149)		1.0–2.4 (W169)		1.0–2.4 (W169)		0.93–2.00 (W149)		0.93–2.00 (W149)	
Secondary bandpass (μm)	0.74–1.0 (Z087)		0.74–1.0 (Z087)		0.74–1.0 (Z087)		0.76–0.98 (Z087)		0.76–0.98 (Z087)	
	W149	Z087	W169	Z087	W169	Z087	W149	Z087	W149	Z087
Zeropoint ^d (mag)	26.315	25.001	26.636	24.922	25.990	24.367	27.554	26.163	27.615	26.387
Exposure time (s)	88	116	85	290	112	412	52	290	46.8	286
Cadence	14.98 min	11.89 hr	14.35 min	12.0 hr	15.0 min	12.0 hr	15.0 min	12.0 hr	15.16 min	12.0 hr
Bias (counts/pix)	380	380	1000	1000	1000	1000	1000	1000	1000	1000
Readout noise ^e (counts/pix)	9.1	9.1	7.6	4.2	9.1	9.1	8.0	8.0	12.12	12.12
Thermal + dark ^f (counts/pix/s)	0.36	0.36	0.76	0.76	0.76	0.76	1.30	0.05	1.072	0.130
Sky background ^g (mag/arcsec ²)	21.48	21.54	21.53	21.48	21.52	21.50	21.47	21.50	21.48	21.55
Sky background (counts/pix/s)	2.78	0.79	3.57	0.77	1.99	0.45	3.28	0.89	3.43	1.04
Error floor (mmag)	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
Saturation ^h (10 ³ counts/pix)	65.5	65.5	80	80	80	80	679	2037	679	679

Notes. Parameters listed in the table are those used in the main simulations whose results are described in Sections 4 and 5 and are not necessarily the same as described in the relevant *WFIRST* reports. Where parameters are incorrect, the impact they would have is judged to be too insignificant to justify a repeated run of the simulations with the correct parameters (see the text for further justification). For correct parameter values the reader should refer to the appropriate Science Definition Team (SDT) report, or the reference information currently listed on the *WFIRST* websites below.

^a https://wfirst.gsfc.nasa.gov/science/WFIRST_Reference_Information.html

^b https://wfirst.ipac.caltech.edu/sims/Param_db.html

^c This accounts for time lost due to constraints on the angle between the Moon and the telescope boresight.

^d Magnitude that produces 1 count per second in the detector.

^e Effective readout noise after multiple non-destructive reads. All values are inaccurate, as they depend on the chosen readout scheme. However, the readout noise will not be larger than the correlated double sampling readout noise of $\sim 20 e^-$, which is still sub-dominant relative to the combination of zodiacal light and blended stars.

^f Sum of dark current and thermal backgrounds (caused by infrared emission of the telescope and its support structures, etc.).

^g Evaluated using zodiacal light model at a season midpoint; in our simulations we use a time-dependent model of the zodiacal background (see Appendix A for details).

^h Effective saturation level after full exposure time. For the designs preceding AFTA, we assumed saturation would occur when the pixel's charge reached the full-well depth. For AFTA we assume that, thanks to multiple reads, useful data can be measured from pixels that saturate after two reads, so for a constant full-well depth, the saturation level increases with exposure time.

fraction of potential general observer observations. These considerations must feed into the spacecraft hardware requirements while simultaneously being constrained by practical design considerations in an iterative process. An example of an observing time line that results from this process is given in the Spergel et al. (2015) report.

While constructing a sample observing schedule for *WFIRST* is a complicated optimization task, requirements set by the nature of microlensing events significantly simplify the process for the microlensing survey. First, the microlensing event rate is highly concentrated toward the Galactic bulge, close to the ecliptic. This means that a spacecraft with a single solar panel structure parallel to its telescope's optical axis can only perform microlensing observations twice per year when the Galactic bulge lies perpendicular to the Sun-spacecraft axis. Second, microlensing events last roughly twice the microlensing

timescale t_E , with $2t_E \sim 60$ days, and planetary deviations last between an hour and a day (see, e.g., Gaudi 2012). This means that a microlensing survey must observe for at least 60 days in order to characterize the whole microlensing event, while also observing at high cadence continually for periods $\gtrsim 1$ day in order to catch and characterize microlensing events. To operate at maximum efficiency it should continuously observe for the entire duration of its survey windows. Finally, the survey requirement to detect ~ 100 Earth-mass planets combined with a detection efficiency of ~ 0.01 per event and a microlensing event rate of around a few $\times 10^{-5} \text{ yr}^{-1} \text{ star}^{-1}$ imposes a requirement of monitoring around a few $\times 10^8$ star years over the duration of the survey.

The duration of planetary deviations places a requirement on the cadence of the microlensing observations. The timescale of planetary deviations is comparable to the Einstein crossing

Figure 1. Total throughput curves for each of the mission designs, compared to spectra of stars of different spectral types, undergoing differing amounts of extinction. The spectrum of a $T_{\text{eff}} = 5800$ K, $\log g = 4.5$ G-dwarf taken from the NEXTGEN grid (Hauschildt et al. 1999) is plotted with no extinction (orange line) and with $A_H = 0.66$ (gray line), which is typical for the expected *WFIRST* fields shown in Section 3.3. G-dwarfs will be the bluest stars that will act as source stars in significant numbers, because more massive stars have evolved off the main sequence in the old bulge population. The y-axis units of the spectra are proportional to the photons per unit wavelength ($dN/d\lambda$), but each is arbitrarily normalized. The throughput curves show the total system throughput including detector quantum efficiency, and are only shown for the Z087 and W149 filters. The *WFIRST* microlensing survey will likely use a wider selection of filters than this.

timescale of an isolated lens of the same mass, $t_E \approx 2$ hr $\sqrt{M/M_\oplus}$, and the deviation must be sampled by several data points in order to robustly extract parameters. Furthermore, finite-source effects, typically lasting ~ 1 hour, carry information about the angular Einstein radius θ_E and need to be resolved by several data points. Combined, these require an observing cadence of ~ 15 min. If stars are well resolved, accurate photometry is possible for much of the bulge main sequence in exposures ~ 1 min on ~ 1 m class telescopes, and so it should be possible to observe between five and 10 fields within the cadence requirements if the observatory can slew fast enough.

The microlensing event rate is highest within a few degrees of the Galactic center, but these regions are also affected by a large amount of extinction. Observations in the near-infrared drastically reduce the effect of the extinction. The *WFIRST* microlensing survey maximizes its photometric precision by using a wide (1–2 μm) bandpass for most of its observations, shown in Figure 1. Combining the wide filter with its wide 0.28 deg^2 field of view, the current design of *WFIRST* can monitor a sufficient number of microlensing events with sufficient precision with ~ 400 days of microlensing observations. Less frequent observations will be taken in more typical broadband filters. Here we assume Z087 in order to measure the colors of microlensing source stars and to measure color-dependent centroid shifts for luminous lenses when the source and lens separate; the range of intrinsic Z087–W149 colors of stars is shown in Figure 22 in Appendix A.1. Note that *WFIRST* magnitudes are on the AB system (Oke & Gunn 1983), and all magnitudes in this paper will be expressed in this system, unless denoted by a subscript Vega.

WFIRST's microlensing survey therefore looks similar across all designs of the spacecraft, with 72 continuous days of observations occurring around vernal and autumnal equinoxes. Six of these seasons are required, with three occurring at the start of the mission and three at the end in order to maximize the baseline over which

Table 2
The *WFIRST* Microlensing Survey at a Glance

Area	1.96 deg^2
Baseline	4.5 yr
Seasons	6×72 days
W149 Exposures	$\sim 41,000$ per field
W149 Cadence	15 min
W149 Saturation	~ 14.8
Phot. Precision	0.01 mag @ W149 ~ 21.15
Z087 Exposures	~ 860 per field
Z087 Saturation	~ 13.9
Z087 Cadence	$\lesssim 12$ hr
Stars (W149 < 15)	$\sim 0.3 \times 10^6$
Stars (W149 < 17)	$\sim 1.4 \times 10^6$
Stars (W149 < 19)	$\sim 5.8 \times 10^6$
Stars (W149 < 21)	$\sim 38 \times 10^6$
Stars (W149 < 23)	$\sim 110 \times 10^6$
Stars (W149 < 25)	$\sim 240 \times 10^6$
Microlensing events $ u_0 < 1$	$\sim 27,000$
Microlensing events $ u_0 < 3$	$\sim 54,000$
Planet detections ($0.1\text{--}10^4 M_\oplus$)	~ 1400
Planet detections ($< 3 M_\oplus$)	~ 200

Note. Assumes the Cycle 7 design. Saturation estimates assumes the brightest pixel accumulates 10^5 electrons before the first read. Star counts have been corrected for the Besançon model's under-prediction (see Section 3.2.1). The exposure time and cadence of observations in the Z087 and other filters has not been set; we have assumed a 12 hr cadence here, but observations in the other filters are likely to be more frequent.

relative lens-source proper motion can be measured (see, e.g., Bennett et al. 2007). The 2.4 m telescope designs of *WFIRST* have a smaller field of view than the ~ 1 m class designs, but can monitor significantly fainter stars at a given photometric precision due to their smaller PSF and larger collecting area, resulting in a similar number of fields being required to reach the same number of stars, despite the difference in field of view. After these two effects cancel, the designs with larger-diameter mirrors come out as significantly more capable scientifically due to the improvement in their ability to measure relative lens-source proper motions. Table 2 summarizes the parameters of the latest iteration of the *WFIRST* microlensing survey design and the survey yields that we will describe in later sections.

3. Simulating the *WFIRST* Microlensing Survey

We performed our simulations using the GULLS code, of which we only give a brief overview here, and refer the reader to Penny et al. (2013, hereafter P13) for full details.⁸ In order to fully simulate *WFIRST* we have made a number of upgrades to GULLS, which are described in Appendix A.

GULLS simulates large numbers of individual microlensing events involving source and lens stars that are drawn from star catalogs produced by a population synthesis Galactic model. Source stars are drawn from a catalog with a faint magnitude limit (here $H_{\text{Vega}} = 25$), and lens stars from a catalog with no magnitude limit; source-lens pairs where the distance of the source is less than the distance of the lens are rejected. Each catalog is drawn from a small solid angle $\delta\Omega$, but represents a larger $0.25^\circ \times 0.25^\circ$ sight line at its specified Galactic coordinates (ℓ , b). The impact parameter u_0 and time of the event t_0 are drawn from uniform distributions with limits

⁸ Note that in P13 the software was called MAB μ LS, but was renamed to disambiguate it from the MAB μ LS online tool (Awiphan et al. 2016).

Figure 2. Left column: section of a Vista Variables in the Via Lactea (VVV) H band image (Saito et al. 2012) from near $(\ell, b) = (1^\circ 1, -1^\circ 2)$, which lies close to the center of the expected *WFIRST* fields. Right four columns: simulated images in the primary wide band of the IDRM, DRM1, DRM2, and AFTA *WFIRST* designs of the same mock star field drawn from the Besançon model sight line at $(\ell, b) = (1^\circ 1, -1^\circ 2)$. The top panels show a 1×1 arcmin² region and the bottom panels show a 4.6×4.6 arcsec² ($\approx 13\times$) zoom-in. The pixel sizes are $0''.339$, $0''.18$, $0''.18$, $0''.18$ and $0''.11$ from left to right respectively. Note that the apparent dark, tenuous, serpentine feature on the left side of the simulated images is a result of random fluctuations in the stellar density, and is not due to spatially varying extinction (e.g., a dust lane). The VVV image is based on data products from observations made with the ESO Telescopes at the La Silla or Paranal Observatories under ESO programme ID 179.B-2002.

$[-u_{0,\max}, +u_{0,\max}]$ and $[0, T_{\text{sim}}]$, respectively, where $u_{0,\max} = 3$ is the maximum impact parameter and T_{sim} is the simulation duration.

Each simulated event i is assigned a normalized weight w_i proportional to its contribution to the total event rate in the sight line

$$w_i = 0.25^2 \text{deg}^2 f_{1106, \text{WFIRST}} \Gamma_{\text{deg}^2} T_{\text{sim}} u_{0,\max} \frac{2\mu_{\text{rel},i} \theta_{\text{E},i}}{W}, \quad (1)$$

where Γ_{deg^2} is the event rate per square degree computed via Monte Carlo integration of the event rate using the source and lens catalogs (see P13 and Awiphan et al. 2016 for details), $f_{1106, \text{WFIRST}}$ is the event rate scaling factor that we use to scale the event rate computed from the Galactic model to match measured event rates (see Section 3.2 for details), $\mu_{\text{rel},i}$ is the relative lens–source proper motion of simulated event i , $\theta_{\text{E},i}$ is the angular Einstein radius of event i , and

$$W = \sum_i 2\mu_{\text{rel},i} \theta_{\text{E},i} \quad (2)$$

is the sum of un-normalized “event rate weights” for all simulated events in a given sight line. As such, the sum of w_i for all events is simply the number of microlensing events we expect to occur in the sight line during the simulation duration with source stars matching the source catalog’s selection criteria. Similarly, the prediction for the number of events matching a given criteria (e.g., a $\Delta\chi^2 > 160$ detection threshold due to a planetary deviation) is simply the sum of normalized weights of events that pass the cut, e.g.,

$$N(\Delta\chi^2 > 160) = \sum_i w_i H(\Delta\chi_i^2 - 160), \quad (3)$$

where $H(x)$ is the Heaviside step function.

Binary (planetary) microlensing light curves are computed using a combination of the hexadecapole approximation (Pejcha & Heyrovský 2009; Gould 2008), contour integration (Gould & Gaucherel 1997; Dominik 1998), and rayshooting (when errors are detected in the contour integration routines; Kayser et al. 1986). Realistic photometry of each event is simulated by constructing images of star fields (drawn from the same

population synthesis Galactic model) for each observatory and filter considered, such that the same stars populate images with different pixel scales and filters. The PSF of the baseline source and lens stars are added at the same position on the image. As the event evolves, the source star brightness is updated and photometry is performed on the image for each data point. For some of the simulations we have implemented a faster photometry scheme that bypasses the need to create a realization of an image for each data point, and which is described in Appendix A. Figure 2 shows examples of simulated images for IDRM, DRM1, DRM2, and AFTA designs compared to a ground-based infrared image, and Figure 3 shows an example simulated color image comparing *WFIRST*’s performance to a simulated ground-based optical telescope in a field typical for the *WFIRST* microlensing survey.

Each star is added using realistic numerical PSFs that are integrated over the detector pixels for a range of sub-pixel offsets and stored in a lookup table for rapid access. For IDRM, DRM1, and DRM2, which all have unobstructed apertures, we used an Airy function averaged over the bandpass of the filter. For AFTA we used numerical PSFs produced using the ZEMAX software package (provided by the *WFIRST* project). For the Z087 filter we used the monochromatic PSF computed at $1 \mu\text{m}$ and for the wide filter we averaged the PSFs computed at 1.0 , 1.5 , and $2.0 \mu\text{m}$ with equal weights. This crude integration procedure insufficiently samples the changing size of the Airy rings as a function of wavelength, so the resulting PSFs have much more prominent higher spatial frequency rings than the actual PSF. The real PSF will be much smoother in the wings (see, e.g., Gould et al. 2014b). The spacing of the unrealistic rings is smaller than the photometric aperture we use, so the inaccurate PSF will have little effect on our results because maxima and minima will average out over the aperture). For the Cycle 7 design we used well sampled numerical PSFs generated using the WEBBPSF tool (Perrin et al. 2012) with parameters from Cycle 5; while the diffraction spikes of these PSFs are rotated 90° relative to the Cycle 7 design, this has no practical effect on our simulated results.

The capabilities of crowded field photometric techniques are approximated by performing aperture photometry on the image with fixed pointing. We found that a 3×3 pixel square aperture produced the best results in the crowded *WFIRST* fields. This simple photometry scheme enables us to accurately simulate all

Figure 3. Simulated color images of an example bulge field at $(\ell, b) = (0^\circ, -1^\circ 5')$ imaged using *WFIRST*’s Cycle 7 detector, compared with a ground-based observatory based on OGLE’s 1.3 m telescope (e.g., Udalski et al. 2015a) in optical filters. The *WFIRST* image is built from a single simulated exposure of 290, 52, and 145 s in Z087, W149, and F184 filters, respectively; the OGLE image is built from single simulated exposures of 150, 125, and 100 s in V, R, and I filters, respectively, i.e., typical of the standard OGLE survey exposures. Note the different sizes of the images compared to the previous figure, and that at least some of the *WFIRST* fields will be amenable to observations with ground-based optical telescopes.

the sources of photon and detector noise that arise in the conversion of photons to data units on an image of minimal size. Aperture photometry is sub-optimal in crowded fields, but we can use this fact to compensate for the effect of any un-modeled causes of additional photometric noise resulting from imperfect data analysis (the precision of any method can only asymptotically approach the theoretically possible photon noise, and sometimes may be far from it), systematic errors (e.g., we do not simulate pointing shifts or variations in pixel response), or data loss. The resultant photometric precision as a function of magnitude is shown in Figure 4, assuming no blending. Properly simulating all sources of systematic or red noise would require a more detailed simulation of the photometry pipeline (for example, by performing difference imaging on images that undergo pointing shifts) and would be significantly more computationally expensive as a significantly larger image would need to be recomputed for each data point. Rather than do this, we simply add in quadrature a Gaussian systematic error floor to the photometry we measure.

We note that we have not correctly simulated the readout schemes employed by the HAWAII HgCdTe detectors *WFIRST* will use. Our simulations simulate the CCD readout process, i.e., an image is exposed for a time t_{exp} before being read out pixel by pixel by one or a small number of amplifiers in a destructive process. Individual pixels in an infrared HgCdTe array have their own amplifier and can be read out non-destructively multiple times per exposure at a chosen rate. HgCdTe amplifiers typically have higher read noise than CCD amplifiers, and so multiple non-destructive reads are employed to reduce the effective read noise in the image. Multiple image “frames” can potentially be stored and downlinked, or processed on-board the spacecraft, enabling retrieval of useful

Figure 4. Single-epoch photometric precision for isolated point sources as a function of magnitude for the Cycle 7 design’s assumed exposure time (46.8 s) assuming no blending. The vertical dashed line indicates the approximate point of saturation in a single read.

data from pixels that would saturate in the full exposure time or that get hit by cosmic rays.

Our photometry simulations assume a single readout at the end of an exposure with a gain of $1 \text{ e}^-/\text{ADU}$. The actual gain value will be different, but any digitization uncertainty will be small compared to the readout noise. We approximate the effective read out noise in *WFIRST* images using the erratum correction of the formula given by Rauscher et al. (2007), based on the correlated double sampling readout noise requirements of each design, an assumed read out rate, and the exposure time. The full-well depth parameter in our simulations is applied to the full exposure time, so we increase the detector full well depth requirement by a factor of $t_{\text{exp}}/t_{\text{read}}$, where t_{read} is the time interval between reads of a given pixel, to simulate the ability to extract a measurement from a pixel that does not saturate before the first read. This workaround results in an underestimate of the Poisson noise component of photometry, but the addition of a 0.001 mag systematic uncertainty in quadrature to the final photometric measurement prevents a severe underestimate of the uncertainty in such situations. The use of only nine pixels in the photometric aperture is also conservative for bright stars. We note that it should be possible to extract accurate photometry from any pixels that do not saturate before the first read (see Gould et al. 2014a).

To assess whether a simulated event contains a detectable planet we use a simple $\Delta\chi^2$ selection criterion:

$$\Delta\chi^2 \equiv \chi_{\text{FSPL}}^2 - \chi_{\text{true}}^2 > 160, \quad (4)$$

where χ_{FSPL} is the χ^2 of the simulated data light curve relative to the best-fitting finite source point lens (FSPL) model light curve (Witt & Mao 1994), and χ_{true}^2 is the χ^2 of the simulated data relative to the true simulated light curve. In practice we only fit an FSPL model if a point-source–point-lens (PSPL) model fit produces a $\Delta\chi^2$ above the detection threshold. We do not consider whether the light curve can be distinguished from potentially ambiguous binary lens models or binary source models.

Our choice of $\Delta\chi^2$ threshold is the de facto standard among microlensing simulations (e.g., Bennett & Rhie 2002; Bennett et al. 2003; Penny et al. 2013; Henderson et al. 2014a). Yee et al. (2012, 2013) discussed the issue of the detection threshold in survey data for high-magnification events, and

concluded that for one particular event a clear planetary anomaly in the full data set might be marginally undetectable in a truncated survey data set at $\Delta\chi^2 \approx 170$. For a uniform survey data set, and a search that included low-magnification events, Suzuki et al. (2016) used a $\Delta\chi^2$ threshold of 100. We expect systematic errors for a space-based survey to be lower than for a ground-based survey, therefore our choice of $\Delta\chi^2 = 160$ should be reasonably conservative. Additionally, except near the edges of its survey sensitivity, the number of planets *WFIRST* can detect is only weakly dependent of $\Delta\chi^2$ as discussed in Section 5. This means that our yields will be relatively insensitive to any inaccuracies in our simulations or models that affect $\Delta\chi^2$. Additionally, because we have chosen relatively conservative assumptions for the systematic noise floor and the $\Delta\chi^2$ threshold, it is possible that the yield of the hardest-to-detect planets could be significantly larger than we predict.

For the smallest-mass planets we do not expect binary lens ambiguity be an issue for many events. Most low-mass planet detections will come from planetary anomalies in the wings of low-magnification events. In such events the caustic location is well constrained and hence also the projected lens–source separation s . Once s is constrained, the caustic size and anomaly duration scales only with the mass ratio q of the lens as $q^{1/2}$ (e.g., Han 2006). Binary source stars with extreme flux ratios can potentially produce false positives for low-mass planetary microlensing (Gaudi et al. 1998). As *WFIRST* will observe source stars much closer to the bottom of the luminosity function, and the near-infrared luminosity function is shallower than the optical luminosity function, we can expect a smaller fraction of *WFIRST*’s binary source stars to have the properties required to mimic a planetary microlensing event. We leave a detailed reassessment of the importance of binary source star false positives for *WFIRST*’s microlensing survey to future work.

3.1. Galactic Model

In this work we use version 1106 of the Besançon Galactic model, hereafter BGM1106. This version of the model is described in full detail by P13 and references therein. It is intermediate to the original, publicly available version (Robin et al. 2003), and a more recent version (Robin et al. 2012). It also differs from the model versions used by Kerins et al. (2009) and Awiphan et al. (2016) to compute maps of microlensing observables.

As the model has been detailed in other papers we only give an overview of the most important features here. The BGM1106 bulge is a boxy triaxial structure following the Dwek et al. (1995) G_2 model with scale lengths of (1.63, 0.51, 0.39) kpc and orientated with the long axis $12^\circ 5'$ from the Sun–Galactic center line. The thin disk uses the Einasto (1979) density law with a scale length of 2.36 kpc for all but the youngest stars, which have a scale length of 5 kpc. The disk has a central hole with a scale length of 1.31 kpc, except for the youngest stars where the hole scale length is 3 kpc. The disk scale height is set by self-consistency requirements between kinematics and Galactic potential (Bienayme et al. 1987). The model also has thick-disk and halo components, but they do not provide a significant fraction of sources or lenses. The full form of the density laws are given in Table 3 of Robin et al. (2003).

Stellar magnitudes are computed from stellar evolution models and model atmospheres based on stellar ages

determined by separate star formation histories and metallicity distributions of the different components. Stellar masses are drawn from an initial mass function (IMF) that differs between the disk and bulge. Each is a broken power law, $dN/dM \propto M^\alpha$, where M is the stellar mass, with $\alpha = -1.6$ for $0.079 < M < 1 M_\odot$ and $\alpha = -3$ for $M > 1 M_\odot$ in the disk, and $\alpha = -1$ for $0.15 < M < 0.7 M_\odot$ and $\alpha = -2.35$ for $M > 0.7 M_\odot$ in the bulge. Extinction is determined using the 3D reddening map of Marshall et al. (2006). This map estimates $E(J - K)$ reddening at various distances and with a resolution of $0^\circ.25 \times 0^\circ.25$ on the sky. Reddening is converted to extinction in each band using the Cardelli et al. (1989) extinction law with a value of total to selective extinction $R_V = 3.1$.

3.2. Normalizing the Event Rate

GULLS computes microlensing event rates by performing Monte Carlo integration of star catalogs produced by the population synthesis Galactic model. In P13 we found that BGM1106 underpredicted the microlensing optical depth by a factor of $f_{\text{od},\text{P13}} = 1.8$ and star counts in Baade’s window by a factor of $f_{\text{sc},\text{P13}} = 1.3$. To account for this we applied a correction factor

$$f_{1106,\text{P13}} = f_{\text{od},\text{P13}} f_{\text{sc},\text{P13}} = 1.8 \times 1.3 = 2.33 \quad (5)$$

to the BGM1106 event rates. Here we will update this event rate correction factor by making comparisons of the Galactic model predictions to new star counts and microlensing event rate measurements.

3.2.1. Comparison to Star Counts

P13 used a comparison between the BGM1106 and *Hubble Space Telescope* (*HST*) star counts in Baade’s window at $(\ell, b) = (1.13, -3.76)$ as measured by Holtzman et al. (1998) to derive a partial correction to the event rate of $f_{\text{sc},\text{P13}} = 1.30$. This field lies more than 2° further away from the Galactic plane than the center of the likely *WFIRST* fields at $b \approx -1^\circ.7$. The Sagittarius Window Eclipsing Extrasolar Planet Search (SWEPS; Sahu et al. 2006) field, originally studied by Kuijken & Rich (2002), lies at $(\ell, b) = (1^\circ.25, -2^\circ.65)$ and is significantly closer to the *WFIRST* fields, but still slightly outside the nominal survey area. The field has been observed by *HST* multiple times over a long time baseline, enabling extremely deep proper motion measurements (Clarkson et al. 2008; Calamida et al. 2015) and now star counts (Calamida et al. 2015). By comparing star counts closer to the *WFIRST* fields we can hope to reduce the impact of any extrapolation errors when estimating an event rate correction.

Calamida et al. (2015) measured the magnitude distribution of bulge stars by selecting stars with a proper motion cut designed to exclude disk stars. Calamida et al. corrected for completeness using artificial star tests, but we add additional corrections for the efficiency of the proper motion cut (34%; A. Calamida 2015, private communication) and the field area ($3/3 \times 3/3$), in order to plot the absolute stellar density as a function of magnitude in Figure 5. We do not consider the bins at the extremes of the magnitude distribution which are likely affected by saturation or large incompleteness. To compare to the observed distribution, we computed the magnitude distribution of bulge stars in the BGM1106 sight

Table 3
Raw Simulation Planet Yields: Log-uniform Mass Function

Mission	IDRM	DRM1	DRM2	AFTA	WFIRST Cycle 7
Duration	432 days	432 days	266 days	357 days	432 days
Area	2.06 deg ²	2.64 deg ²	3.52 deg ²	2.82 deg ²	1.97 deg ²
Rate Norm.	2.81	2.81	2.81	2.81	2.81
Mass ($M_{\oplus}$)					
0.1	8.1 ± 0.6	11.1 ± 0.4	7.0 ± 0.2	18.0 ± 0.3	9.9 ± 0.4
1	79.8 ± 3.6	87.5 ± 2.1	66.1 ± 1.5	138 ± 2	87.5 ± 2.6
10	366 ± 15	496 ± 9	350 ± 6	643 ± 6	439 ± 8
100	1610 ± 47	2110 ± 39	1500 ± 26	2440 ± 51	1780 ± 84
1000	5480 ± 150	6610 ± 110	4790 ± 80	7670 ± 130	5210 ± 86
10 ⁴	12700 ± 230	15400 ± 190	11400 ± 130	17500 ± 200	11300 ± 150

Note. The table presents planet yields for each simulated survey, with the uncertainties due to Poisson shot noise in the drawn event parameters; this uncertainty does not include any systematic component due, e.g., to the normalization of event rates; see Section 6 for a discussion of the magnitude of these errors. The survey duration, total field area, and event rate normalization are shown in the header lines of the table. The yields assume a planet of fixed mass and an occurrence rate of one planet per decade of semimajor axis in the range $0.3 \leq a < 30$ au per star.

Figure 5. Comparison of Besançon model star counts for the bulge population as a function of magnitude $F814W_{\text{Vega}}$ at $(\ell, b) = (1^\circ 35', -2^\circ 70')$ to those measured in the *HST* SWEEPS field at $(\ell, b) = (1^\circ 25', -2^\circ 65')$, which lies close to the expected *WFIRST* fields. Green squares show bulge-only star counts from *HST* (Calamida et al. 2015) and diamonds show counts of red giant branch stars in the same area from OGLE-III (Szymański et al. 2011). The *HST* stars were selected to be bulge stars by proper motion cuts, and have been corrected for the approximate efficiency of this cut. The solid black line shows the BGM1106 prediction, with error bars denoting the Poisson uncertainty of the catalogs. While there are differences in the detailed shape of the star count distribution, integrated over the range $F814W = 19\text{--}26.5$, BGM1106 underpredicts the total number of stars by 33%; the blue line shows BGM1106 scaled up by this factor. The dashed brown line shows the BGM1106 model star counts if the mass function is changed in the bulge to match mass function 1 of Sumi et al. (2011), namely a broken power law with slopes of -1.3 and -2.0 (dN/dM) each side of a break at $0.7 M_{\odot}$.

line at $(\ell, b) = (1^\circ 35', -2^\circ 70')$, which is the closest to the SWEEPS field.

BGM1106 matches the measured magnitude distribution reasonably well between $F814W_{\text{Vega}} = 19.5$ and 23, though with minor differences in shape. BGM1106 starts to significantly underpredict the number of stars fainter than $F814W_{\text{Vega}} = 23$. Brighter than $F814W_{\text{Vega}} = 22.9$, Calamida et al. (2015) find 10% more stars than BGM1106 predicts. Integrated over the magnitude range $F814W_{\text{Vega}} = 19\text{--}26.5$ BGM1106 underpredicts star counts by 33%. The magnitude of the discrepancy is very similar to that we found between the BGM1106 and the Holtzman et al. (1998) luminosity function, giving us some confidence that there is no significant gradient

in BGM1106’s star count discrepancy. We adopt the star count scaling factor of $f_{\text{sc}} = 1.33$.

The cause of the discrepancy between model and data can be partially explained by BGM1106’s choice of IMF in the bulge, $dN/dM \propto M^{-1.0}$. Adopting a more reasonable mass function (e.g., mass function number 1 from Sumi et al. 2011, $dN/dM \propto M^{-2.0}$ for $M > 0.7 M_{\odot}$ and $dN/dM \propto M^{-1.3}$ for $0.08 < M < 0.7 M_{\odot}$), and assuming that the BGM1106 star counts were normalized using turn-off stars of $1.0 M_{\odot}$, produces the luminosity function prediction shown by the dashed line in the plot. This mass function overpredicts the star counts fainter than $F814W_{\text{Vega}} \approx 20$, but better matches the shape of the entire observed luminosity function between $I_{\text{Vega}} \approx 19\text{--}26$. Both the original and modified BGM1106 mass functions slightly underpredict the number of giant branch star counts from the same sight line detected by OGLE (Szymański et al. 2011), which have not been corrected for incompleteness. We note here that we do not adopt an alternative mass function (e.g., mass function 1 from Sumi et al. 2011), but discuss the impact of the mass function on our results in Section 6.

3.2.2. Comparison to Microlensing Event Rates

Since writing P13, Sumi et al. (2013) published measurements of the microlensing event rate toward the bulge, in addition to optical depth measurements. Measurements of the event rate per source star allow a more direct route to estimating any corrections to the model’s predicted event rates, so here we only perform a comparison to the event rates and not the optical depths. For the comparison we use the event rates from Sumi & Penny (2016), which corrected the Sumi et al. (2013) event rates and optical depths for a systematic error in estimates of the number of source stars monitored.

Sumi et al. (2013) present event rates for two samples of events, the “extended red clump” (ERC) sample composed of events with source stars brighter than $I = 17.5$ and colors selected to only include the bulge giant branch, and the “all stars” (AS) sample composed of all events with $I < 20$ and no color cut. We selected star catalogs from BGM1106 to match these samples, and computed event rates per source Γ by Monte Carlo integration over this source catalog and a lens catalog with no magnitude or color cuts (see Awiphan et al. 2016 for a detailed description of such calculations). The small angle of the Galactic bar to the line of sight in BGM1106 ($\sim 12^\circ$) results in the bulk of our AS sample source stars lying

Figure 6. Comparison of the microlensing event rate per source predicted by the Besançon model to the Sumi & Penny (2016) revision of measurements by Sumi et al. (2013). Black data points show measurements for all source stars, while red data points show measurements for the extended red clump source stars (see the text for details). The thin line shows BGM1106’s prediction of extended red clump event rates, and the thick red line shows this prediction after multiplication by the best-fit scaling parameter $f_{\Gamma} = 2.11 \pm 0.29$.

in front of most of the bulge stars. This leads to significantly smaller event rates per source than would be expected for a more reasonable bar angle of $\sim 30^\circ$ (e.g., Cao et al. 2013; Wegg & Gerhard 2013), which could lead us to overcorrecting the event rates. We therefore only compare to the ERC sample event rates.

Figure 6 shows the predicted model event rates, averaged over the range $-0.53 < \ell \leq +2.73$, and the data from Sumi & Penny (2016), which were averaged over the range $|\ell| < 5$. BGM1106 predicts a lower event rate than is measured. At latitudes $|b| < 1.5$ the predicted ERC event rate turns over, likely because extinction begins to limit the range of distances over which significant numbers of bulge giants pass the ERC color and magnitude cuts; at more negative latitudes, where the observations we compare to were made, the extinction likely has a smaller impact. We find that multiplying the BGM1106 event rates by a constant scaling factor $f_{\Gamma} = 2.11 \pm 0.29$ yields a good match to the observed ERC rates, with $\chi^2 = 1.58$ for 3 degrees of freedom. Although the model predictions cover a smaller range of ℓ than the measurements, Sumi & Penny (2016) results binned by ℓ indicate only a relatively weak dependence of Γ on $|\ell|$.

3.2.3. Adopted Event Rate Scaling

In P13 we scaled the microlensing event rates computed using BGM1106 by making the assumption that all of the relevant distributions (e.g., kinematics, mass, and density) were reasonable, but that there could be errors in the normalization of the numbers of source and lens stars. To make a correction for the number of source stars we directly compared the BGM1106 predictions to deep star counts measured by the *HST*. To estimate the correction for the number of lens stars, we compared model predictions to measurements of the microlensing optical depth. This has the advantage that, should the density distribution and mass function of stars in the model be reasonable, the necessary correction to the event rate due to lenses should scale with the optical depth discrepancy between model and data. However, as we have described in the

preceding subsections, and will expand on in Section 6, the density distribution (specifically the bar angle of the bulge), the bulge mass function, and the kinematics of bulge stars in BGM1106 are inconsistent with current measurements. These will affect the event rate and optical depth in different ways that are not trivial to calculate. This makes any simple scaling of the event rate based on optical depth comparisons suspect. In contrast, a scaling based on measured event rates is far more direct with fewer assumptions. Therefore, to correct the event rates predicted by the Galactic model, we adopt the event rate scaling factor

$$f_{1106,WFIRST} = f_{sc} f_{\Gamma}, \quad (6)$$

where

$$f_{sc} = 1.33, \text{ and } f_{\Gamma} = 2.11, \quad (7)$$

for a total event rate correction of

$$f_{1106,WFIRST} = 2.81, \quad (8)$$

which is about 20% larger than the scaling adopted in P13. We will discuss the impact of uncertainties and inaccuracies in the Galactic model beyond the event rate scaling in Section 6.

We apply the $f_{1106,WFIRST}$ scaling throughout the paper as our fiducial event rate normalization. However, we will also present our main results with the scalings used for each of the *WFIRST* reports in order to aid comparison to these earlier works; these results using obsolete scalings are presented in Table 5 in Section 4. We note, however, that when applying the obsolete scaling, we did not include a factor of 1.475 in the scaling that was used in Green et al. (2012). This factor was used to account for a factor of 2.2 discrepancy in the microlensing detection efficiency of our GULLS simulations and simulations performed by D. Bennett, based on Bennett & Rhie (2002) and updated for simulations of *WFIRST* (Green et al. 2011); 1.475 was the geometric mean of the relative detection efficiencies for planets of $1 M_{\oplus}$ with a period of 2 yr. The cause of the difference in detection efficiencies was not conclusively tracked down. However, at fixed period, the projected separation $s \propto M^{-5/6}$, and the detection efficiency is a strong function of s , so a difference between the host mass function of the simulations (see Sections 3.2.1 and 6.2.3) is likely to cause a significant difference in the detection efficiency at fixed period. Averaged over a range of semimajor axis, as we have done in the simulations presented in Section 4, we can expect any difference in the detection efficiency at fixed planet mass to be significantly smaller.

3.3. The WFIRST Fields

For the IDRM, DRM1, DRM2, and AFTA simulations, the field placement was not rigorously optimized. We show the fields we adopted for each design in Figure 7. For IDRM, DRM1, and DRM2, the field placement is significantly different from what it would be in reality if each design were flown. This is due to uncertainties in the orientation of the detectors in the instrument bay. We therefore chose the simplest field orientation we could, aligning the principle axes of the fields with Galactic latitude and longitude. Note, however, that this is an optimistic assumption, as the extinction and event rate at zeroth order depend strongly on b but weakly on ℓ , so detector orientations that align the long axis with ℓ are likely to be close to optimal. For the IDRM, DRM1, and DRM2 field layouts we accounted for gaps between detectors

Figure 7. Assumed field placement for each *WFIRST* design, plotted over a map of *H*-band extinction (Gonzalez et al. 2012). The gaps between IDRM, DRM1, and DRM2 fields were included to mimic the effects of gaps between detectors. For the AFTA and Cycle 7 simulations we accounted for the individual detector placement within each field more carefully, so fields are close-butted (note the curved focal plane). Note also that in reality the 1 m class designs would also likely have curved detector layouts and that, unlike the AFTA and Cycle 7 designs, the fields would probably not be orientated with their principle axes aligned with Galactic coordinates. The black diamond in the top panels shows the location of the *HST* SWEEPS field. In the Cycle 7 panel, colored dots show the detection rate of $1 M_{\oplus}$ planets per square degree as a function of position for the Cycle 7 design. A version of the Cycle 7 plot is available at <https://github.com/mtpenny/wfirst-ml-figures>.

in the focal plan by placing twice the sum of all chip gaps between each of the fields.

For AFTA we considered the field layout more carefully. The telescope instrument bay already exists, setting the orientation of the field and constraining the layout of detectors within it. Coincidentally the orientation of the WFI focal plane for AFTA is within a couple of degrees of alignment with Galactic coordinates. From spring to fall seasons the orientation of the detector will be rotated by 180° , which means that with the curved geometry of the active focal plane, the fields observed will not be exactly the same. For the layout shown this results in $\sim 90\%$ of stars that fall on a chip in spring seasons also falling on a chip in the fall seasons. Occasional

Figure 8. Number of detections as a function of planet mass for each *WFIRST* design and for two different mass functions. Also shown is the distribution of *Kepler* candidates assuming a mass-radius relation of $(M/M_{\oplus}) = (R/R_{\oplus})^{2.06}$ (Lissauer et al. 2011). Letters indicate the masses of solar system planets. The lower panel shows the yields of the DRM designs relative to that of Cycle 7.

gap filling dithers could be used to ensure some observations in both spring and fall for all events.

Between the AFTA design and Cycle 7, the *WFIRST* WFI was redesigned and consequently the field orientation was rotated by 90° . Additionally, the spacecraft's slew and settle time estimates were updated, and were more than twice that we had assumed for the AFTA design. These two changes led us to conduct an optimization of the field layouts. This optimization is described in Section 5.4. With this field layout we can expect to detect $\sim 27,000$ microlensing events with $|u_0| < 1$ and roughly twice this with $|u_0| < 3$ during the course of the mission. While there are three times as many events with $|u_0| < 3$ compared to $|u_0| < 1$, the maximum magnification of a Paczynski (1986) single-lens light curve at $|u_0| = 3$ is only 1.017, compared to 1.34 at $|u_0| = 1$, so only on brighter stars will it be possible for *WFIRST* to detect these low-magnification events.

4. Baseline *WFIRST* Planet Yields

To assess the performance of each *WFIRST* design we ran a series of simulations to investigate the number of planets that would be detected during the *WFIRST* microlensing survey. To assess the performance of the survey over a broad range of masses and orbits we simulated single-planet events with fixed masses in the range $0.1 \leq M \leq 10^4 M_{\oplus}$ with semimajor axes distributed logarithmically in the range $0.3 \leq a < 30$ au—roughly a factor of 10 either side of the typical Einstein radius (2–3 au).

Table 3 shows the raw results of these simulations for each *WFIRST* design, and the same numbers are plotted in Figure 8. The table is not particularly useful for assessing the number of planets that *WFIRST* will detect, because the mass function that it implies (one planet per decade of mass and semimajor axis; we will call this the log-uniform mass function) is a significant overestimate at larger masses. It is, however, the most convenient form from which other mass functions can be applied. The results for IDRM1, DRM1, and DRM2 were first presented in the [DRM report](#), though using a different set of event rate corrections (see Section 3.2 for a description of the event rate corrections) and included a small number of

Table 4
Best-estimate Planet Yields: Fiducial Mass Function

Mission	IDRM	DRM1	DRM2	AFTA	<i>WFIRST</i> Cycle 7
Duration	432 days	432 days	266 days	357 days	432 days
Area	2.06 deg ²	2.64 deg ²	3.52 deg ²	2.82 deg ²	1.97 deg²
Rate Norm.	2.81	2.81	2.81	2.81	2.81
Mass ($M_{\oplus}$)					
0.1	16.7 ± 1.3	22.8 ± 0.8	14.4 ± 0.5	37.1 ± 0.6	20.5 ± 0.8
1	164 ± 8	180 ± 4	136 ± 3	284 ± 3	180 ± 5
10	455 ± 19	615 ± 11	434 ± 7	799 ± 8	545 ± 9
100	371 ± 11	488 ± 9	346 ± 6	563 ± 12	412 ± 19
1000	236 ± 6	284 ± 5	206 ± 4	330 ± 6	224 ± 4
10⁴	101 ± 2	124 ± 2	91.3 ± 1.1	141 ± 2	90.7 ± 1.2
Total (0.1–10⁴ $M_{\oplus}$)	1294	1653	1183	2084	1428

Note. In this table the yields presented in Table 3 have been multiplied by our fiducial mass function (Equation (9)). The total survey yield is found by integration of the tabulated values with the extended trapezoidal rule.

additional detections of planets with semimajor axis between $0.03 \leq a < 0.3$ au. The results for DRM2 differ further from the report due to a correction factor used to the correction of a mistake in the number of fields and the exposure times that was made in the simulations presented in the Green et al. (2012) report. The AFTA simulations were rerun completely to incorporate the changes introduced at Cycle 4 of the *WFIRST* design. In hindsight these yields are overly optimistic due to an unrealistic assumption of the spacecraft’s slew performance. The yields for Cycle 7 stand as the most realistic and up-to-date estimates, and so we have made these bold in the table.

Table 4 presents our fiducial estimate for the planet yield of *WFIRST*. To compute these yields we multiply the raw yields from Table 3 by the following fiducial form of the planet mass function:

$$\frac{d^2N}{d \log M_p d \log a} = \begin{cases} 0.24 \text{ dex}^{-2} \left(\frac{M_p}{95 M_{\oplus}} \right)^{-0.73} & \text{if } M_p \geq 5.2 M_{\oplus}, \\ 2 \text{ dex}^{-2} & \text{if } M_p < 5.2 M_{\oplus}, \end{cases} \quad (9)$$

which we will refer to hereafter as the fiducial mass function. This mass function is based on the power-law bound planet mass function measured by Cassan et al. (2012) from microlensing observations, but saturated at a value of two planets per star below a mass of $5.2 M_{\oplus}$, which is roughly where Cassan et al. lost sufficient sensitivity to measure the mass function. For context, this saturation value is comparable to the planet density of the inner solar system between ~ 0.2 and 2 au. We discuss the impact on our results of more recent determinations of the mass ratio function in Section 6. To estimate the total number of planets *WFIRST* will detect we integrated the resulting numbers using the trapezoidal rule; for the Cycle 7 design this results in an expected total yield of ~ 1400 planets, noting that it does not include free-floating planets which would add a few hundred to this figure (Spergel et al. 2015). This yield is smaller than *Kepler*’s total yield, but of the same order of magnitude, suggesting that a similar precision in demographics will be achievable. We note that in order to aid comparison to previous *WFIRST* SDT reports

(Green et al. 2012; Spergel et al. 2015), in Table 5 we have provided a version of Table 4 that uses the event rate scaling factors that were used in the respective reports (see Section 3.2 for details).

Figure 8 shows the simulation yields of Tables 3 and 4 graphically, and compares them to *Kepler*. To estimate the *Kepler* mass function we applied the simplistic mass–radius $M_p/M_{\oplus} = (R_p/R_{\oplus})^{2.06}$ relation of Lissauer et al. (2011) to the *Kepler* candidates with $\text{koi_score} > 0.5$.⁹ The first point to note is the large number of planets that *WFIRST* is expected to detect, using any of the designs and across a wide range of masses. As we discuss later, the cold planet mass function below $\sim 10 M_{\oplus}$ is almost completely unconstrained at present, meaning that *WFIRST* will add at least two orders of magnitude in mass sensitivity beyond current knowledge.

The lower panel of Figure 8 compares the relative yields of the previous *WFIRST* designs to the current Cycle 7 design as a function of planet mass. Except for the highest-mass planets, the Cycle 7 design outperforms the IDRM and DRM2 designs in terms of planet detection yield, though note that DRM2 had a significantly shorter total survey duration. The DRM1 design has a $\sim 10\%$ higher yield than the current Cycle 7 design at low masses, but a larger difference at large planet masses. The comparison of only the planet detection yields between Cycle 7 and the ~ 1 m class designs is, however, unfair. Cycle 7’s 2.4 m mirror enables factor of ~ 2 improvements in the measurement of host and planet masses through lens flux, image elongation, and color-dependent centroid shifts (Bennett et al. 2007), which the planet detection yield alone does not account for. We will discuss these measurements in Section 4.2.

We have not shown the AFTA yields in Figure 8 because we consider them to be unrealistic. The significant drop in yield from the AFTA design to Cycle 7 deserves some discussion, however. While there were many changes between the designs, one change dominates the reduction in yield. This is the adoption of more accurate estimates of the slew and settle time of the spacecraft, and is not a result of descope of mission hardware. For the AFTA designs, we had assumed values for slew performance that were the same as for the smaller IDRM and DRM designs, which was unrealistic given the larger mass and moment of inertia associated with a larger mirror, secondary mirror support structure, and spacecraft bus. Our

⁹ <https://exoplanetarchive.ipac.caltech.edu/>

Figure 9. Comparison of the *WFIRST* Cycle 7 design sensitivity to that of *Kepler* in the planet mass–semimajor axis plane. The red line shows an approximation of the *Kepler* planet detection limit based on Burke et al. (2015). Blue shading shows the number of *WFIRST* planet detections during the mission if there is one planet per star at a given mass and semimajor axis point; this is directly proportional to the average detection efficiency. The thick, dark blue line is an functional fit to the three-detection per mission contour, while the lighter blue line barely visible beneath it is the actual contour. Red dots show *Kepler* candidate and confirmed planets; black dots show all other known planets extracted from the NASA exoplanet archive (accessed 2018 February 28 Akeson et al. 2013). Blue dots show a simulated realization of the planets detected by the *WFIRST* microlensing survey, assuming our fiducial planet mass function (Equation (9)), though note that in constructing this sample of simulated detections we did not simulate planets smaller than $0.03 M_{\oplus}$ or with semimajor axis less than 0.3 au. Solar system bodies are shown by their images, including the satellites Ganymede, Titan, and the Moon at the semimajor axis of their hosts. Images of the solar system planets credit to NASA. The data and scripts used to make this plot are available at <https://github.com/mtpenny/wfirst-ml-figures>.

Table 5
Obsolete Planet Yields from *WFIRST*-AFTA Final Report: Fiducial Mass Function

Mission	IDRM	DRM1	DRM2	AFTA	<i>WFIRST</i> Cycle 7
Duration	432 days	432 days	266 days	357 days	432 days
Area	2.06 deg ²	2.64 deg ²	3.52 deg ²	2.82 deg ²	1.97 deg ²
Rate Norm.	2.41×1.475	2.46×1.475	2.42×1.475	2.46×1.475	2.46×1.475
Mass ($M_{\oplus}$)					
0.1	21.1 ± 1.6	29.4 ± 1.0	18.3 ± 0.6	47.8 ± 0.8	25.9 ± 1.0
1	208 ± 10	232 ± 6	173 ± 4	367 ± 4	228 ± 7
10	575 ± 24	793 ± 14	551 ± 9	1030 ± 10	690 ± 12
100	470 ± 14	629 ± 12	439 ± 8	726 ± 15	522 ± 25
1000	298 ± 8	367 ± 6	261 ± 4	426 ± 7	283 ± 5
10 ⁴	128 ± 2	160 ± 2	116 ± 1	181 ± 2	115 ± 2
Total (0.1–10 ⁴ $M_{\oplus}$)	1636	2131	1499	2687	1806

Note. As Table 4, but using the event rate (~ 2.4) and detection efficiency compromise (1.475) scalings that were used in the *WFIRST*-AFTA final report (Spergel et al. 2015). These yields should be considered obsolete, and are only presented to aid comparison with the previous *WFIRST* reports.

AFTA results are further unrealistic, because we applied the ~ 0.4 slew time to all slews between fields, when some slews will be longer. These optimistic assumptions for the previous designs, if corrected, would likely result in a less optimal field layout and a reduction in yields, though perhaps relatively smaller than the drop from AFTA to Cycle 7. If it were possible to use attitude control systems that provide significantly faster

slew performance (e.g., control moment gyros instead of reaction wheels), then significant gains in the yield of the *WFIRST* microlensing survey could be realized.

4.1. Sensitivity to Moon-mass Objects

To trace out the approximate limits of the sensitivity of *WFIRST* to low-mass planets, as well as wide- and close-separation planets,

$$M = 2.02 M_{\text{Moon}} \quad a = 5.20 \text{ AU} \quad M_{\star} = 0.29 M_{\odot} \quad \Delta\chi^2 = 710$$

Figure 10. Example light curve of a $0.025 M_{\oplus}$ (Ganymede-mass) bound planet detection from simulations of the AFTA design. Black and red data points are in the W149 and Z087 filters, respectively. The blue curve shows the underlying “true” light curve and the green line shows the best-fit single-lens light curve.

we ran simulations on a grid in planet mass and semimajor axis over the ranges $-2 \leq \log(M_p/M_{\oplus}) < 4$ and $-2 \leq \log(a/\text{au}) < 2$. We required that the events have impact parameters $|u_0| < 3$ and times of closest approach in the range $0 \leq t_0 < 2011$ days. The full details of the computations are described in Appendix B, including details of how false-positive detections due to numerical errors are accounted for. We note that the simulations were conducted using the parameters of the AFTA design, but we used the analytic estimate method described in Section 5.1 to predict the yield of the Cycle 7 design from the AFTA design simulations. Other than the change in field layout, there is only a small change in exposure time for the analytic approximation to account for, so we expect the uncertainty due to this conversion to be small.

Figure 9 shows the results of the grid computation. The shading shows the planet detection rate (in units of planets per full survey) at each mass–semimajor axis point. The blue sensitivity curve plots the specific point in the mass–semimajor axis plane where three planet detections can be expected in the course of the mission if there is one such planet per star. The sensitivity curve is also a line of constant detection efficiency (e.g., Peale 1997). We found that the sensitivity curve is very well approximated by an analytic function

$$\log(M_p/M_{\oplus}) = \alpha + \beta \log(a/\text{au}) + \gamma \sqrt{\delta^2 + [\log(a/\text{au}) - \epsilon]^2}, \quad (10)$$

where a is the semimajor axis and the parameters take the following values: $\alpha = -3.90$, $\beta = -1.15$, $\gamma = 3.56$, $\delta = 0.783$, $\epsilon = 0.356$. The analytic function (bright blue) is plotted above the sensitivity curve data (pale blue) in the figure, but the sensitivity data curve is almost invisible underneath the analytic fit. It is interesting to note that seven of the eight solar system planets fall within the *WFIRST* sensitivity curve, with only Mercury outside. However, in place of Mercury, when we place the solar system’s moons on the diagram at the orbital separations of their host planets, the Galilean moon Ganymede

lies just within the sensitivity curve. Figure 10 shows an example light curve of such a Ganymede-mass exoplanet that *WFIRST* could detect. If the solar system moons were placed at the minimum point of the analytic curve ($a = 4.2 \text{ au}$, $M = 0.021 M_{\oplus}$), Ganymede ($0.025 M_{\oplus}$), and Titan ($0.023 M_{\oplus}$) would be above the curve and Callisto ($0.018 M_{\oplus}$), Io ($0.015 M_{\oplus}$), and the Moon ($0.012 M_{\oplus}$) would be below it; all other known solar system bodies have masses less than $0.01 M_{\oplus}$. At $1 M_{\oplus}$, the sensitivity curve stretches from 0.5 to 70 au. Removing the constraint that the impact parameter of the host star’s microlensing event be $|u_0| < 3$ results in there being no upper limit on the semimajor axis at which AFTA has sensitivity to Earth-mass planets. However, most of these more distant planets would be seen as free-floating planet candidates due to undetectable magnification from the source star (see Henderson & Shvartzvald 2016 for a discussion of constraining the presence of host stars in free-floating planet candidate events).

The shading in Figure 9 does not accurately represent the distribution of planets that we can expect to detect with *WFIRST*, only *WFIRST*’s sensitivity to them. *WFIRST* has a high-detection efficiency for large-mass planets, but these have consistently been shown to be rare (e.g., Cumming et al. 2008; Howard et al. 2010; Cassan et al. 2012; Fressin et al. 2013; Shvartzvald et al. 2016; Suzuki et al. 2016). To give a better idea of the distribution of planet detections *WFIRST* will detect, we simulated a survey with planets populated from our fiducial mass function. A realization from this simulation is shown as the blue dots in Figure 9, and Figure 11 shows an unrepresentative sample of some events in this realization. Note that the lower-mass limit for this simulation was $M_p = 0.03 M_{\oplus}$ and the smallest semimajor axis was 0.3 au, so a small number of the most extreme planets that could be detected will be missing from this realization. In adding these simulated planets to the plot, we do not account for the uncertainty in measurements of either M_p or a . We expect the majority of *WFIRST*’s planets to have mass measurements from either lens-detection measurements or parallax so, given the large range of masses covered, uncertainties in M_p would not result in a significant change of appearance, and features in the planet mass function should be easily distinguishable. With a mass measurement also comes a measurement of the physical Einstein radius, and so the physical projected separation of the planet and star. This must be deprojected to the actual semimajor axis, which will result in a substantial uncertainty in the estimated value of a relative to the range of a , though for some planets orbital motion measurements can better constrain a (e.g., Bennett et al. 2010b; Penny et al. 2011). Our simulations do not include an eccentricity distribution, so we cannot estimate the uncertainty associated with this deprojection.

4.2. Properties of the WFIRST Microlensing Events

The *WFIRST* microlensing survey will search for microlensing events from fainter sources than are observed from the ground, and its resolution will be sufficient to at least partially resolve lenses and sources over the course of the five year mission, so it is important to consider the properties of the sources, lenses, and blending. Figure 12 plots the distribution of source magnitudes and blending of the microlensing events that the Cycle 7 design will be able to observe. In both plots we show the distribution for every microlensing event that occurs in the *WFIRST* field on sources brighter than $H_{\text{Vega}} = 25$, regardless of whether it will be detected or not (labeled “All μL ”), the events that cause a

Figure 11. Examples of simulated *WFIRST* light curves using the Cycle 7 design parameters, chosen to display the some of the variety of light-curve features that *WFIRST* will detect, or challenges that will impact the analysis of events. The examples demonstrate a light curve with missing peak data due to *WFIRST*'s limited observing seasons (top left), an Earth-mass planet orbiting a $0.15 M_{\odot}$ star at 1 au (top right), a wide orbit planet with a very low amplitude host microlensing event (middle left), an event with a very faint source and high blending (middle right), a high-S/N detection of a massive planet (bottom left), and a low-S/N detection of a massive planet on a wide orbit (bottom right).

microlensing event that is detectable as a $\Delta\chi^2 > 500$ deviation from a flat baseline (labeled “Detected μL ”), and the events with detectable Earth-mass planets (labeled “Detected planets”).

While the number of all microlensing events per magnitude keeps rising beyond a magnitude of $W149 = 25$, the number of

detectable microlensing events exhibits a broad peak between $W149 = 22$ and 24 , before beginning to fall. For planet detections, the source magnitude distribution peaks between $W149 \approx 20$ – 22 , but only begins to fall rapidly fainter than $W149 \approx 24$.

Figure 12. Top: distribution of source magnitudes for different subsets of the simulated microlensing events. “All μL ” is the distribution of all simulated microlensing events, “Detected μL ” is the distribution of events with flux variation detected at $\Delta\chi^2 > 500$ above a flat baseline, and “Detected planets” is the distribution for events with $1 M_{\oplus}$ planet detections. Each distribution is on its own arbitrary scale. Bottom: distribution of the fraction of baseline flux contributed by the source, f_s . Again, each distribution has been arbitrarily scaled.

The lower panel of Figure 12 shows the distribution of blending. The blending parameter f_s is the ratio of source flux to total flux in the photometric aperture when the source is unmagnified. We have measured f_s in the same 3×3 pixel aperture we have used for photometry. Given that the input source magnitude distribution continues to rise toward faint magnitudes, the majority of microlensing events we simulated were significantly blended (i.e., $f_s < 0.5$), despite *WFIRST*’s small PSF. This is also the case for microlensing events that *WFIRST* will detect, though the distribution $dN/d \log f_s$ does peak above $f_s = 0.5$. For events with planets detected, the distribution of f_s is much more skewed toward small amounts of blending $f_s \sim 1$, but there remains a significant tail with large amounts of blending.

Figure 13 shows the relative contribution of the lens star to the total blend flux for detected planets. The lens flux dominates the blended light in fewer than 20% of events with planet detections. However, the majority events will have a lens flux within a factor of 10 of the total blend flux. Without knowing the nature of the blended light (i.e., whether it is due to the PSF wings of bright stars or nearby fainter stars), it is

Figure 13. Cumulative distribution of the contribution of the lens flux to the total blend flux in a 3×3 pixel aperture around events with detected planets. The vertical line shows the point at which the lens provides more than 50% of the flux.

Figure 14. Plot of the lens–source relative proper motion plotted against the lens–source brightness contrast for the simulated planet sample from Figure 9. Horizontal lines are spaced vertically by the proper motion required for the separation to change by $\sim 0.1 \times$ the PSF FWHM in 4.5 yr, the spacing between the first and last microlensing season. Histograms on each axis show the marginalized distributions. Note that μ_{rel} has not been corrected for BGM1106’s high proper motion dispersions as discussed in Section 6.2.2.

difficult to say more about how this added confusion will affect host mass measurements via lens detection (see, e.g., Bennett et al. 2007, 2015; Henderson et al. 2014b).

Figure 14 shows the joint distribution of the lens–source brightness contrast and the lens–source relative proper motion for the simulated sample shown in Figure 9. These are the two intrinsic properties of the event that have the largest impact on whether *WFIRST* will be able to characterize the lens through detection of lens flux and motion relative to the source. The proper motions are not corrected in any way for the overly large proper motion dispersions in BGM1106’s bulge stars, which cause μ_{rel} to be $\sim 15\%$ – 25% too large (see Section 6.2.2 for a detailed discussion). Regardless of this, it can be seen that the majority of source–lens pairs will separate by $\sim 10\%$ of the FWHM of the PSF in the course of the mission, and few cases will separate by more than half the PSF FWHM. By carefully modeling the PSF and the motion of the source and lens, it will

Figure 15. Cumulative distribution of source diameter crossing times $2t_*$ for events with $1 M_{\oplus}$ planet detections and high-magnification events with $|u_0| < 0.05$. Vertical lines indicate multiples of the 15 min observing cadence. Most events will have at least one measurement per source diameter crossing time.

be possible to measure lens fluxes and lens–source separations when the separation is significantly smaller than the PSF (see, e.g., Bennett et al. 2007, 2015; Henderson et al. 2014b). We leave it to future work to simulate these measurements.

The final property of *WFIRST* events we shall examine is the source diameter crossing time $2t_*$, where t_* is the time taken for the source (relative to the lens) to traverse its radius. This is important to consider because, in addition to the requirement of sampling the planetary perturbation, the cadence must be chosen to also sample any finite source effects, which can be used to measure the angular Einstein radius (e.g., Nemiroff & Wickramasinghe 1994; Yoo et al. 2004). With *WFIRST* being sensitive to fainter, and thus smaller, source stars than any previous microlensing survey, there is a possibility that assumptions about the required cadence overlooked this fact. This is especially important to consider, as Chung et al. (2017) have identified a degeneracy in measuring $2t_*$ for some single-lens events when only a single measurement is affected by finite-source effects. Figure 15 shows the cumulative distribution of $2t_*$ for events with $1 M_{\oplus}$ planet detections, and for high-magnification events with $|u_0| < 0.05$, compared to the cadence of the *WFIRST* microlensing survey observations. Over 80% of events with planet detections will have at least two points per $2t_*$, and over 70% will have three points per $2t_*$. This is likely an underestimate because we expect $2t_*$ to be longer than simulated due to the overestimated velocity dispersion of BGM1106’s bulge. High-magnification events detected by *WFIRST* will tend to have shorter $2t_*$ because they can be detected on fainter source stars on average. Even so, $\sim 60\%$ of events will be guaranteed two observations per $2t_*$ and more than this can expect two measurements more often than not in any given $2t_*$ -long time interval. Free-floating planet events will have brighter sources than high-magnification stellar microlensing events on average, so we conclude that for most events 15 min cadence is sufficient to avoid the possible source radius measurement degeneracy (identified by Chung et al. 2017) when there are only a small number of photometric measurements over the part of the light curve affected by finite-source effects.

Figure 16. Cumulative distribution of $\Delta\chi^2$ for different planet masses (solid black lines) ranging from 0.1 to $1000 M_{\oplus}$ in factor-of-ten steps. Gray dashed lines show the power-law fits to the cumulative distributions over a factor of two above and below the adopted $\Delta\chi^2 = 160$ threshold (as indicated by the vertical dashed lines). The slopes of the fits are listed in Table 6.

5. Evaluating the Effect of Changes to Mission Design on Planet Yields

The design of any space mission must balance capabilities with cost. *WFIRST* straddles the boundary between a targeted, single-goal mission for which a focused set of hardware can be optimized within a relatively constrained parameter space, and a general purpose observatory where the breadth of capabilities should be optimized. Combined, *WFIRST*’s primary missions present a relatively broad scope for optimization, though the synergies between the observational requirements of each survey and current economic considerations constrain this scope considerably. In this section we consider the effect of changes in the design of the spacecraft and the survey it carries out on the overall planet yields. We begin by outlining how changes can be quickly estimated analytically in Section 5.1, present the results of two trade study simulations that we use to test the analytic estimates in Sections 5.2 and 5.3, and then apply the analytic estimates to optimize the field choice for the *WFIRST* Cycle 7 design in Section 5.4.

5.1. Analytic Estimates of the Change in Yield

It is possible to estimate the effect of a change in the design of hardware or survey strategy on the total planet yield without performing a full simulation. The only detection criteria we use for bound planets is a cut on $\Delta\chi^2$. Therefore, the distribution of $\Delta\chi^2$ combined with a model of how $\Delta\chi^2$ changes with design can be used to estimate a change in yield. The cumulative distribution of the number of planet detections with $\Delta\chi^2$ greater than a threshold X , $N(\Delta\chi^2 > X)$, can be approximated locally by a power law (e.g., Bennett & Rhie 2002)

$$N(\Delta\chi^2 > X) \propto X^{-\alpha}, \quad (11)$$

as can be seen in Figure 16; the fitted slopes are listed in Table 6. The slope of the power law α is a function of both the planet mass and semimajor axis of the planetary companion. The range of validity of the approximation can extend by more than an order of magnitude in $\Delta\chi^2$ in some cases, though for planets close to the edges of *WFIRST*’s parameter space it becomes increasingly inaccurate. Finally, if one knows the ratio

Table 6
Power-law Slopes Fitted to the Cumulative $\Delta\chi^2$ Distributions

Mass ($M_\oplus$)	α				
	IDRM	DRM1	DRM2	AFTA	Cycle 7
0.1	0.674	0.473	0.520	0.513	0.534
1	0.420	0.364	0.355	0.366	0.399
10	0.324	0.290	0.296	0.310	0.313
100	0.315	0.246	0.241	0.265	0.245
1000	0.268	0.212	0.212	0.223	0.227
10000	0.201	0.168	0.151	0.193	0.204

of the new $\Delta\chi^2$ to the old

$$\delta = \Delta\chi_{\text{new}}^2 / \Delta\chi_{\text{old}}^2, \quad (12)$$

then from Equation (11) the estimate of the yield for the new design N_{new} is simply

$$N_{\text{new}} \approx N_{\text{old}} \delta^\alpha. \quad (13)$$

To increase the range of validity of the approximation a higher-order polynomial could be fit to the local cumulative $\Delta\chi^2$ distribution in log–log space, or the full $\Delta\chi^2$ distribution could be used to directly evaluate the change in yield that corresponds to a given change in $\Delta\chi^2$; we assess these options in Section 5.4.

We can calculate δ for two scenarios by computing the S/N for each scenario, because if the signal is constant $\Delta\chi^2 \propto (\text{S/N})^2$. The S/N of each photometric measurement is

$$\text{S/N} = \frac{N_s}{\sqrt{N_s + N_{\text{BG}} + \sigma_{\text{det}}^2 + \sigma_{\text{sys}}^2}}, \quad (14)$$

where N_s is the number of photons from the source, N_{BG} is the number of detected photons from blended stars and smooth backgrounds, σ_{det} is the uncertainty on the total number of counts due to readout noise, dark current, thermal photons from the spacecraft and other detector effects, and σ_{sys} is a systematic error component, which in this paper we assume is a small constant (0.001) multiplied by N_s . Each element of the S/N will scale differently as parameters of the telescope or survey change, and in general will be a function of magnification that will be different for each microlensing event. It will therefore be difficult to estimate the average ratio of $\Delta\chi^2$ because S/N will not be a simple function of the changing parameters. However, each of the N terms in the equation scale linearly with the photon collection efficacy of each scenario, so if both of the σ^2 components are small and there is nothing to change the ratio of source to background between the different scenarios, S/N will scale as the square root of photon collection efficacy and so δ will scale linearly. This will be the case if the only the exposure time were to change, but not if the mirror diameter were changed, because in the latter case the mirror diameter affects the resolution and hence the ratio of source to background photons will change from event to event.

The concept can be illustrated by an example. We may be interested in how the planet yield would be affected if the observation cadence was halved, keeping all else fixed for this example. On average, halving the number of data points will halve the $\Delta\chi^2$. However, halving the cadence allows the

exposure time to be increased by a factor of $(2t_{\text{exp}} + t_{\text{thead}})/t_{\text{exp}} = 2.73$ for AFTA which, assuming that systematic errors and detector noise are negligible, increases the per-exposure signal to noise by $\sqrt{2.73} = 1.65$. Overall then, the $\Delta\chi^2$ will increase by a factor of $\delta = 1.65^2/2 = 1.37$. Without worrying about the normalization of the power law in Equation (11), we can estimate that, relative to the fiducial case, halving the cadence will result in a yield that is a factor $N_{\text{cadence}/2}/N_{\text{cadence}} \approx 1.37^\alpha$ times larger than the fiducial yield. So, for the examples shown in Figure 16, halving the cadence would result in an increase in yields of $100\% \times (1.37^\alpha - 1) \approx 9\%$ for $100 M_\oplus$ -mass planets where $\alpha = 0.270$. For Earth-mass planets where $\alpha = 0.379$, the increase would be 13%.

Note, however, that the above example also demonstrates a need for caution when applying scalings based on $\Delta\chi^2$. In our pursuit of larger $\Delta\chi^2$, we have neglected the role of sampling. In order for each planet detection to be useful for demographic studies, we not only need to detect a deviation from a single-lens microlensing light curve at a specified significance, we also need to be able to fit the light curve with a unique planetary model and be sure that the deviation is not caused by systematics in the data. By halving the number of data points over a potentially short-lived planetary deviation, we have significantly degraded our power to reject systematic errors and to constrain our light-curve model. These effects could ultimately reduce the number of useful detections by a factor larger than the increase in detections due to the improvement in the average $\Delta\chi^2$. We must therefore take care not to over-interpret any of the results in this section. To properly assess the impact of changes in design therefore requires even more detailed simulations than we have conducted here, which assess not only the detection significance but also the level to which events can be characterized. However, with sensible restrictions on the survey design in place to ensure that detected events will be well characterized (such as restrictions on the minimum cadence), these simulations and analytic estimates can provide useful insight into the effect of design trade-offs.

5.2. Testing Analytic Estimates: Bandpass

The filter bandpass primarily effects the amount of light that reaches the detector, but also influences the width of the PSF. Therefore, the actual effect of the bandpass on the photometric precision will change from event to event due to the differing colors of the source, lens, background, and any blended stars, as well as differing amounts of reddening.

WFIRST's HgCdTe detectors can be designed to have a specified red cutoff wavelength, with options in the range of 2.0–2.4 microns considered for *WFIRST*. Longer cutoff wavelengths allow for larger total throughputs in the wide microlensing band, as well as a longer-wavelength baseline for its more standard filters (Green et al. 2012). However, this added capability comes at the cost of additional spacecraft and instrument cooling, which is needed to reduce the thermal background emitted by the mirrors and other components in the optical path.

To test the impact of bandpass on the planet yield of the *WFIRST* microlensing survey, we ran simulations of DRM1 with 2.0 and 2.4 μm cutoff detectors, using the W149 and W169 filters, respectively. This scenario is also ideal for quantitatively testing the validity of our analytic relative yield estimates. It should be possible to approximate the effect of the bandpass to zeroth order by just considering the total photon

Table 7
Relative Yield of $1 M_{\oplus}$ Planets for DRM1 Comparing W149 and W169 Bandpasses

Simulation Name	W149	W169'	W169
Zeropoint (mag)	26.377	26.377	26.636
Filter	W149	W169	W169
Simulation	0.910	0.915	1
Prediction	0.901	0.917	...

Note. Yields are given relative to the standard W169 simulation. The W169' simulation uses the W169 filter, but a total throughput (zeropoint) equal to the W149 simulation. The filter determines the PSF, brightness of stars relative to the zeropoint magnitude, and the surface brightness of the zodiacal light.

throughput for a source with a spectrum that is representative of sources that will be observed. To test this assumption, we ran an additional simulation that used the W169 filter, but the W149 zeropoint magnitude, which we will refer to as W169'.

The relevant parameters and relative yields of these simulations are listed in Table 7. All other parameters between the three simulations are identical. We note that we should have reduced the amount of thermal noise in the W149 simulation, but neglected to do so. Fortunately, this mistake has essentially no effect on the result because blended light and sky background dominate over the detector noise for detector noise levels this low. The W149 bandpass results in a drop in yield of 9.0% relative to the W169 bandpass. Because each simulation used the same set of simulated events, the statistical uncertainty in the relative yields is significantly smaller than the 1% Poisson uncertainty on the total yield.

We predicted the relative yield using the analytic formalism from Section 5.1. If we ignore detector and systematic noise sources, and assume that both blend stars and source stars have the same color, then the ratio of $\Delta\chi^2$ between W149 and W169 will depend only on the difference in zeropoints and the assumed W149 – W169 color, which we take to be $W149 - W169 = 0.052$, which is the median of source colors for the planets detected in the W169 simulation. Therefore, $\log \delta = -0.4(0.259 + 0.052) = -0.12$, and the predicted reduction in yield for W149 relative to W169 is $(\delta^\alpha - 1) = 9.9\%$ for the $1 M_{\oplus}$ DRM1 value of $\alpha = 0.364$, which is in reasonable agreement with the simulation's value of 9.0%. This estimate does not account for the reduced blending in W149 thanks to the narrower PSF, which might explain why the analytic estimate predicts a slightly larger drop in yield than the simulation produces. The same calculation for the W169' simulation, which only changes the zeropoint of the W169 simulation, predicts a drop in yield of 8.3% compared to the actual simulation result of 8.5%, and is thus a much closer match. The W169' simulation shows that the majority of the change in yields is due to the change in zeropoint associated with a change in bandpass. However, because of the relatively shallow slope of the $\Delta\chi^2$ distribution, even a relatively significant change in bandpass ($\sim 40\%$) results in a significantly smaller change in planet yield ($\sim 10\%$), at least for Earth-mass planets. While we have only simulated the impact for Earth-mass planets, having validated the analytic approximation of the yield change, we can apply it to other planet masses. For example, we predict that switching to a $2.0 \mu\text{m}$ cutoff for DRM1 would reduce $0.1 M_{\oplus}$ yields by 13%.

Figure 17. Relative yield of $1 M_{\oplus}$ planets in the presence of elevated smooth backgrounds. Data points are the results of our simulation and the solid black line is our analytic prediction adopting values of the average surface brightness of blend stars $W149_{\text{blend}} = 19.2 \text{ mag arcsec}^{-2}$ and the average source star brightness $W149_s = 22$. Red and blue dashed and dotted lines show the impact of changing the values of these parameters by $\pm 1 \text{ mag arcsec}^{-2}$ or $\pm 1 \text{ mag}$, respectively.

5.3. Testing Analytic Estimates: Background Light

While there is nothing that one can do to control the amount of astrophysical background light that enters the telescope from diffuse backgrounds, it is nevertheless important to consider the effect that varying levels of background have on the yield of the survey. To investigate the impact of the zodiacal light—the dominant smooth background—and its variations over the course of the *WFIRST* survey, we implemented a time- and position-dependent model of the zodiacal light, which is described in Appendix A.2. We found the impact of variations in the zodiacal light over the course of *WFIRST*'s 72 day seasons to be negligible on the overall yields. We also investigated the impact of adding additional smooth backgrounds to the images. This can be used to estimate the impact of observing when the moon lies near to the microlensing fields (if *WFIRST* is in a geosynchronous Earth orbit, as was baselined for early versions of the 2.4 m design). Figure 17 shows the yield of Earth-mass planets in 0.3–30 au orbits as a function of the surface brightness of the added background relative to the case with just zodiacal light (which is $W149 \approx 21.5 \text{ mag arcsec}^{-2}$). The yield drops steadily as the background surface brightness increases, but the drop is not severe unless the additional background is very bright. In order to cut the yields in half, the additional background must exceed the zodiacal light by a factor of ~ 80 . This test is also relevant to increased thermal noise backgrounds for differing telescope and instrument operating temperatures, and the choice of whether to include a cold pupil mask on the W149 filter.

To test the analytic estimate, we assumed that a typical source will have a magnitude of $W149_s = 22$ (the peak of the detected microlensing event source magnitude distribution; see Figure 12) and that the influence of blended stars would be a smooth background of $W149_{\text{blend}} = 19.2 \text{ mag arcsec}^{-2}$, based on summing up the flux of stars in BGM1106 with magnitudes $H_{\text{vega}} > 15$ (which was arbitrarily chosen to be the boundary of one of our BGM1106 catalogs; see Table 1 of P13). Each simulated scenario used the same set of microlensing events, meaning that the uncertainty on the relative yield was significantly smaller than the Poisson uncertainty on any of

the individual simulations. To compute the uncertainty on the relative yields, we split the simulated sample into 10 parts and computed the variance in relative yield measured for each of the subsamples.

There is good qualitative agreement between the analytic estimate and the simulation results but, moving from right to left in the plot, the relative detection rate for analytic estimate falls off less quickly than the simulations until providing a better match at backgrounds brighter than $W149 = 18$ mag arcsec $^{-2}$. Additionally, our analytic estimate depends on two model parameters (the typical source magnitude and the effective surface brightness of unresolved and blended stars). We appear to have chosen their values well, but the choices were somewhat arbitrary, and others could have been justified. This demonstrates that care needs to be taken when using analytic estimates without simulations to anchor them.

5.4. Applying Analytic Estimates: Optimizing Field Selection for Cycle 7

More accurate models of the mass, inertia and reaction wheel complement of the spacecraft led to significantly longer estimates of the slew and settle time for the spacecraft in Cycle 7. This prompted us to perform a more detailed accounting of the survey's overheads and a re-optimization of the number and placement of the survey's fields. This process relied on analytic yield change estimates to quickly assess the yield for a large number of potential exposure times as described below. To simplify the optimization process, we constrained the cadence to be fixed at 15 min, and optimized only for the yield of Earth-mass planets in a broad range of orbits (i.e., with a logarithmically distributed between 0.3 and 30 au).

Before beginning the optimization process, we took heed of our above warning to treat the analytic yield change estimates carefully. To test the accuracy of the analytic yield estimates, we simulated identical surveys using the Cycle 7 mission parameters at four different exposure times spanning the expected range of values to be seriously considered in the optimization exercise: 26, 52, 104, and 130 s. We used the $\Delta\chi^2$ distributions of the 130 s and 52 s simulations to predict the yield of each other simulation. We used both a linear (i.e., power law), and a quadratic fit to $\log N(\Delta\chi^2 > X)$ versus $\log \Delta\chi^2$, as well as a direct evaluation of the cumulative $\Delta\chi^2$ distribution. The results are shown in Figure 18. No particular method of estimating the change in yield showed consistently better accuracy. The error in the approximation grew by $\sim 5\%$ for every factor of two change in exposure time, whether the 52 s or 130 s simulations were used as the basis of the approximation. This level of inaccuracy is a reasonable price to pay for the computational cost savings the method provides.

Unlike with the previous designs, we evaluated the cumulative slew time for a given set of fields using estimates of the cumulative slew, settle and detector reset time as a function of slew angle provided by the *WFIRST* project office. This calculation was done for a large number of candidate field layouts with varying geometries and numbers of fields, with the best path being selected by a brute force solution of the traveling salesman problem. The exposure time was divided evenly between the number of fields from the remainder of 15 min cadence minus the total slew time. We constructed a map of the $1 M_{\oplus}$ planet detection rate per unit area for each

Figure 18. Fractional error of analytic yield estimates compared to actual simulated yields for simulations with different exposure times, and all else held fixed (including cadence). Analytic estimates of the yield were computed based on the cumulative $\Delta\chi^2$ distributions of simulations run using 52 s (open circles) and 130 s (filled circles) exposure times for the main W149 survey observations. The $\Delta\chi^2$ distribution was modeled by linear fits to the cumulative $\Delta\chi^2$ distribution, quadratic fits to the distribution, and direct evaluation of the distribution (indicated by point color).

candidate layout's exposure time t_{exp} by direct evaluation of the 52 s simulation's cumulative $\Delta\chi^2$ distribution at $X = 52 \text{ s}/t_{\text{exp}}$ individually for each sight line in the map. The total number of detections for a each layout was then evaluated using polygon clipping (Murta 2015) to estimate the fraction of the area represented by each sight line that falls within a given chip of each field.

The results of the optimization exercise are shown in the top panel of Figure 19 for three different slew time versus slew angle profiles: the first is a constant slew, settle, and reset time of 38 s independent of slew distance (this was assumed in the AFTA simulations), and the other two profiles are for *WFIRST* Cycle 7 with all of its reaction wheels operational and with one reaction wheel inoperational. For each number of fields, we consider several possible layouts. For the AFTA slew times, the optimum number of fields was 10 or larger (we did not consider layouts with more than 10 fields). For more realistic Cycle 7 slew times there is a broad optimum of between five and eight fields. We adopt a slightly sub-optimal field layout of seven fields (shown in Figure 7) to allow for some margin in yields. Note that the relatively coarse resolution of our event rate map, uncertainties in the yields for each sight line, and use of the analytic yield estimates prevent accurate determination of the true optimum field layout within the broad optimum we find. With one reaction wheel inoperational the optimum number of fields would be fewer, at five or six. This optimum is somewhat sharper than the optimum for all wheels operational. The factor of ~ 2.2 reduction in estimated slew performance between AFTA and Cycle 7 results in a large reduction in planet yield of $\sim 50\%$, which accounts for essentially all of the difference between the AFTA and Cycle 7 yields that we find. The optimum number of fields will depend on planet mass, as shown in the lower panel of Figure 19, which shows the yield as a function of field number for 0.1, 10, and 1000 $M_{\oplus}$ planets.

6. Discussion

The statistical power of an exoplanet survey to infer demographics is directly related to the expected yield of the

Figure 19. Top: $1 M_{\oplus}$ planet yield as a function of the number of fields for simulations of the Cycle 7 design, the Cycle 7 design with a failed reaction wheel, and the overly optimistic AFTA slow times. Several possible layouts are considered for each number of fields. Bottom: planet yield relative to the adopted field layout (see Figure 7) as a function of the number of fields for the nominal Cycle 7 slow times and for 0.1, 10, and 1000 $M_{\oplus}$ planets. Different numbers of fields would maximize the yield for different planet masses.

survey assuming a given exoplanet population. The ability to accurately estimate the survey yields is therefore an important input into mission design. Nevertheless, yield predictions have numerous potential sources of uncertainty, and it is just as important to understand these. We therefore devote this section to summarizing the sources of uncertainty in our yield predictions and suggesting ways in which this uncertainty can be reduced.

The sources of uncertainty in our results can be broken down into three broad categories. The first is due to our ability to simulate how the spacecraft collects data and how it will be processed. The second is due to our ability to measure and model the astrophysical components that produce microlensing events, i.e., the Galaxy and its stellar populations. Finally, our assumptions of the planetary population impacts the mission yields that we predict.

6.1. Simulation Uncertainties

A principal concern when building a simulation is the balance between realism and computational cost. One is invariably forced to make compromises on the former in order to obtain a manageable run time. By building simulated images of the *WFIRST* microlensing fields, combined with a model of

smooth backgrounds, our simulations should reasonably capture all significant sources of photon noise. Our simulation of the detectors is somewhat simplistic; we assume that *WFIRST*'s HAWAII-4RG detectors behave like CCDs, which is probably reasonable given that photon noise always dominates over read noise for the microlensing fields and their exposure times. The most significant form of uncertainty in the data collection and processing category is the processing element. Our simulations perform a simple aperture photometry noise calculation and limit its precision with a constant 1 mmag term added in quadrature. Ultimately, the use of a small, fixed, unweighted aperture should result in an underestimate of the achievable photon-noise-limited photometric precision, which helps to offset our inability to simulate all of the imperfections between photons entering the telescope and a photometric measurement. However, there are a number of additional steps between measuring photometry and declaring a planet detection that could be significantly affected by sources of systematic noise (instrumental and/or astrophysical) that we do not simulate. Estimating the impact of these sources of noise is challenging, and would likely require a full end-to-end simulation. Realistic noise simulations based on lab tests of H4RG detectors (e.g., Rauscher 2015) will help in this regard. We note that the impact of any change in the simulation of photometry, insofar as it changes the photometric precision by a uniform scaling factor, can be estimated by changing the $\Delta\chi^2$ threshold for declaring a detection.

Astrophysical variability is a potentially important source of systematic photometry errors and false positives. For the brightest of *WFIRST*'s main sequence source stars, with G and K spectral types, where photometric systematics may dominate over photon noise, Gilliland et al. (2011) have shown using *Kepler* data that the majority of stars have variability on ~ 6 hr timescales of < 0.05 mmag, so a factor of 20 smaller than our adopted systematic error component. On longer timescales the median variability is ~ 0.1 mmag (Gilliland et al. 2015). In *WFIRST*'s bandpass the variability amplitudes would be smaller than this because stellar surface features caused by temperature variations have smaller contrasts in the infrared than the optical. For M dwarfs, Goulding et al. (2012) found that fewer than 1% of M dwarfs had periodic variability with $> 0.9\%$ amplitude in the *J* band. Stellar flares can resemble caustic entrances, but have decay timescales that are at least a factor of a few shorter than microlensing source crossing timescales, and infrared amplitudes far smaller than in the optical (e.g., Tofflemire et al. 2012), so are not expected to be a significant source of false-positive planet detections. Even if flares are not mischaracterized as planets, they could destroy the ability to detect planets when they occur. However, Davenport et al. (2012) estimate the rate of flares larger than 0.01 mag in the *J* band to be in the range of $0.01\text{--}0.001 \text{ hr}^{-1}$ for spectral types M0–M5, so flares should not significantly impact *WFIRST*'s planet finding duty cycle for early M-dwarf sources. Later M dwarfs have significantly higher activity levels and flare rates (e.g., Davenport et al. 2012; Schmidt et al. 2015), but will not make up a significant fraction of the *WFIRST* source stars. So we do not expect any form of astrophysical variability to significantly affect our yield estimates.

Another important uncertainty in the simulations is the ability to convert a detection, i.e., a signal above the $\Delta\chi^2$ threshold, into a bone fide planet with measured parameters.

This process can be affected by various discrete and continuous degeneracies (see Gaudi 2012 for a review) that can lead to ambiguity between planetary interpretations and stellar binary lenses and sources as the cause of the light-curve anomaly. These ambiguous events can be dealt with by Bayesian probability accounting, but naturally add uncertainty to any inferences, especially if they constitute a significant fraction of the potential planet sample. Reassuringly, Suzuki et al. (2016) found only one out of 23 events in their systematically selected sample of planets had an ambiguous binary lens interpretation, and six out of 23 events suffered a close-wide degeneracy that impacted the measurement of the projected separation s , but did not significantly affect the mass ratio q . The improved photometry possible from space will likely resolve some fraction of the degeneracies and ambiguities that are seen in ground-based data, but the more subtle features detectable in space-based data may introduce a higher fraction of ambiguous events; e.g., in a simulation of a high-cadence microlensing survey with uniform photometry, Zhu et al. (2014) found 55% of planets were detected without caustic crossings. While we cannot say for certain how many planet detections will be badly affected by degeneracies, we expect it to be a small fraction of events with detected anomalies.

In this work we have not simulated the measurement of planet and host masses that observations from space enable. The majority of such measurements will be made by detection of lens light in the high-resolution *WFIRST* images, either as an elongation of the PSF or a color-dependent centroid shift (Bennett et al. 2007). For the MPF mission design, with pixels more than twice the size of *WFIRST*'s Cycle 7 design, Bennett et al. (2010b) estimate that more than half of the planetary events will have better than 10% mass measurements via some form of direct detection of lens light. Bhattacharya et al. (2018) recently compared the precision estimates of Bennett et al. (2007, 2010b) to measurements they made using *HST* and found that the precision estimates were reasonably accurate. Therefore, we expect most of the planets *WFIRST* detects will also have their host and planet masses measured via direct detection of lens light in the *WFIRST* survey images. Note, however, that these estimates do not account for the potential contamination of the measurement by bound stellar companions either to the source or lens (Henderson 2015; Koshimoto et al. 2017).

6.2. Galactic Model Uncertainties

The microlensing event rate and the properties of the microlensing events will depend on the distribution of stars in the Galaxy, their kinematics, and their masses. No model of the Galaxy will be able to fully capture its complexities, and so our estimates will have some degree of uncertainty due to any shortcomings of the model. In Section 3.2 we estimated the corrections necessary to match the microlensing event rate of the model. In this section we examine further the possible causes for the BGM1106's underprediction of event rates and the impact of any model uncertainties or errors on the properties of the microlensing events. We note that this discussion only applies specifically to version 1106 of the Besançon model that we have used in this paper.

6.2.1. Bar Angle

In BGM1106 the Galactic bulge is modeled as a triaxial bar with an angle of 12.5° to the Sun–Galactic center line and a major axis scale length of 1.63 kpc. This is in contrast to modeling of the distribution of red clump giants found in the OGLE and Vista Variables in the Via Lactea (VVV) surveys, which find a bar angle $\sim 30^\circ$ and a major axis scale length of ~ 0.7 kpc (Stanek et al. 1994; Rattenbury et al. 2007; Cao et al. 2013; Wegg & Gerhard 2013). With no distance information, both models can reproduce the 2D distribution of bulge stars on the sky (note that $1.63 \text{ kpc} \sin 12.5^\circ \approx 0.7 \text{ kpc} \sin 30^\circ \approx 0.35 \text{ kpc}$). Red clump stars are standard candles (Stanek et al. 1994), and so can trace out the third dimension of the bulge density distribution. Along the line of sight $(\ell, b) = (1.25, -2.65)$ BGM1106 predicts a distance modulus dispersion of 0.36 mag for bulge stars, which is much larger than the value of 0.20 mag that Nataf et al. (2013) measured from OGLE clump giants after subtracting in quadrature an intrinsic magnitude dispersion of 0.09 mag and an extinction dispersion of ~ 0.11 mag (the total observed magnitude dispersion of the red clump is therefore 0.24 mag). Simion et al. (2017), working with the Galaxia code that implements the 2003 Besançon model (Sharma et al. 2011), found that slightly smaller bar angles of 20° – 25° provided the best fit to VVV red clump counts, but that there was some degeneracy between the bar angle and the red clump dispersion due to sources other than distance dispersion.

We found in Section 3.2.1 that BGM1106 only slightly underpredicts bulge star counts, so the principal impact of the bar angle is only to spread bulge stars along the line of sight. This was confirmed by comparing the red clump star counts of Nataf et al. (2013) as a function of (ℓ, b) to the total stellar mass of the BGM1106 bulge.

The line-of-sight distribution of lenses and sources affects the distribution of microlensing event properties. The Einstein radius depends on the relative distances of the source and lens as

$$r_E = \sqrt{\frac{4G}{c^2} MD_1(1 - D_1/D_s)}, \quad (15)$$

where G and c are the gravitational constant and speed of light, respectively, and D_1 and D_s are the lens and source distances, respectively. The larger line-of-sight distance dispersion for the BGM1106 relative to that measured will result in $\sim 15\%$ larger Einstein radii for events with bulge lenses and bulge sources (bulge–bulge lensing), and a smaller impact on events with disk lenses and bulge sources (bulge–disk lensing). Event time-scales will be larger by the same degree, and the ratio of bulge to disk lenses will also be increased. However, with only a maximum effect of $\sim 15\%$, the impact will be relatively minor.

6.2.2. Bulge Kinematics

In addition to the bar angle, the kinematics of the bulge stars are not in agreement with measurements. We compared the predicted BGM1106 proper motions to *HST* proper motion measurements in the $\sim 11 \text{ arcmin}^2$ SWEEPS field (Clarkson et al. 2008). Clarkson et al. separated bulge and disk populations by selecting stars above the bulge main sequence turn-off where the disk's main-sequence stars were well

separated from the giant branch, which is dominated by bulge stars. For this comparison, we combined the BGM1106 catalogs of the two sight lines closest to the SWEEPS field at $(\ell, b) = (1^\circ.1, -2^\circ.7)$ and $(1^\circ.35, -2^\circ.7)$ to improve the statistics for these bright stars. We roughly mimicked the Clarkson et al. (2008) selection in our BGM1106 catalogs by selecting stars with $15.95 < I_{AB} < 17.95$ and assigned those with $(I - J)_{AB} < 0.5$ to the blue (disk proxy) population and those with $(I - J)_{AB} > 0.58$ to the red (bulge proxy) population. The two catalogs combined represent stars drawn from a solid angle of 1.44 arcmin^2 , and there are a total of 37 stars in the blue disk proxy sample and 105 stars in the red bulge proxy sample.

Clarkson et al. (2008) measured their proper motions in an arbitrary reference frame, so we can only compare the blue–red proper motion offsets and the proper motion dispersions. They found an offset between the blue and red population proper motions of $(\Delta\mu_\ell, \Delta\mu_b) = (3.24 \pm 0.15, -0.81 \pm 0.12) \text{ mas yr}^{-1}$, where the Δ represents blue minus red. For the BGM1106 proper motions we find $(\Delta\mu_\ell, \Delta\mu_b) = (3.53 \pm 0.65, -0.12 \pm 0.32) \text{ mas yr}^{-1}$, which are largely consistent with each other.

Clarkson et al. do not report the proper motion dispersions they measure, but we are able to extract them from their Figure 21, finding $(\sigma_\ell, \sigma_b) = (2.2, 1.3) \text{ mas yr}^{-1}$ for the blue (disk) population and $(\sigma_\ell, \sigma_b) = (3.0, 2.8) \text{ mas yr}^{-1}$ for the red (bulge) population. Individual proper motion uncertainties in the *HST* data are likely below 0.3 mas yr^{-1} for each star, so have a negligible impact on the measured dispersions, and the sample size is larger than our comparison sample so the statistical uncertainty in the estimates of the *HST* proper motion dispersions will be insignificant in our comparison. For the blue BGM1106 population we find proper motion dispersions of $(\sigma_\ell, \sigma_b) = (2.47 \pm 0.29, 1.11 \pm 0.13) \text{ mas yr}^{-1}$, largely consistent with the *HST* measurements in both axes. For the red BGM1106 population we find $(\sigma_\ell, \sigma_b) = (5.19 \pm 0.36, 2.64 \pm 0.18) \text{ mas yr}^{-1}$; the latitudinal dispersion is consistent with the *HST* measurements, but the longitudinal dispersion is too large by a factor of 1.73 ± 0.12 .

We can get an idea for the cause of the discrepancy between BGM1106 and the data by looking at the proper motion vector point diagram and the longitudinal proper motion plotted against distance, both shown in Figure 20. The first thing to notice is that the color–magnitude selection does a good job of separating near-side disk stars from bulge stars, admittedly with a small degree of cross-contamination between the disk and bulge populations. Also, the selected stars do not appear to be a significantly biased subset of the underlying population. This means that we cannot ascribe the discrepancy to a difference in the selection of the stars for each population between the observation and the model. The random velocity dispersions in the BGM1106 bulge population are $(\sigma_U, \sigma_V, \sigma_W) = (113, 115, 100) \text{ km s}^{-1}$, which correspond to proper motion dispersions of the order of 3 mas yr^{-1} at a distance of 8 kpc, i.e., enough to account for all of the measured *HST* value of σ_ℓ . In addition to the dispersion component, the BGM1106 bulge stars also have an additional solid body rotation component, rotating at $40 \text{ km s}^{-1} \text{ kpc}^{-1}$. The combination of the longer bar and its small angle lead to a range of solid body rotation velocities from $\sim -60 \text{ km s}^{-1}$ on the near end of the bar to $\sim +50 \text{ km s}^{-1}$ as the sight line leaves the far side. This results in a range of -1.9 to $+1.1 \text{ mas yr}^{-1}$ in proper motion,

Figure 20. Comparison of BGM1106 model proper motions with those measured using *HST* by Clarkson et al. (2008). Top: proper motion vector point diagram. Small points are all potential source stars ($H_{\text{vega}} < 25$) from BGM1106 (light blue are disk stars, while light red are bulge stars). Larger blue and red points are BGM1106 stars belonging to disk and bulge proxy populations, respectively, with selections designed to replicate those of Clarkson et al. Light blue and red lines are the 1σ proper motion dispersion contours measured by Clarkson et al. for disk and bulge proxy stars, respectively, while the darker lines are the same 1σ dispersion contours, but for BGM1106 disk and bulge proxy stars. Bottom: distance versus ℓ proper motion for the same BGM1106 stars as in the top panel. The CMD cuts do a good job of isolating disk and bulge populations, though with some cross-contamination. The BGM1106 disk proper motions match the data well, but the bulge proper motions have a dispersion that is too large in the ℓ direction. The dotted line shows the distance of the Galactic center in the model.

which would correspond to an additional dispersion of $\sim 0.87 \text{ mas yr}^{-1}$ to be added in quadrature. Finally, inspection of BGM1106’s V -component velocities (in a UVW system) as a function of distance suggests that, in addition the solid body rotation and the random dispersions, there is an additional component similar to the rotation curve of a stellar disk in a dark matter halo. This results in a $\sim 300 \text{ km s}^{-1}$ offset between the mean velocities of bulge stars that are $\sim 1 \text{ kpc}$ from the Galactic center on opposite sides. This causes a rapid change of $\sim 8 \text{ mas yr}^{-1}$ in the mean proper motion of stars as function of distance at the distance of the Galactic center. This can be seen in the lower panel of Figure 20 as an offset in proper motion at a distance of $\sim 8 \text{ kpc}$. This offset results in additional $\sim 4 \text{ mas yr}^{-1}$ to be added in quadrature to the proper motion dispersion. Each of the three sources of dispersion combined in quadrature results

in a dispersion of 5.1 mas yr^{-1} , which is consistent with the value of $\sigma_\ell = 5.19 \pm 0.36 \text{ mas yr}^{-1}$ we measured from the catalogs.

The addition of a potential quasi-circular velocity component to the bulge stars' velocities appears to be in error, because we expect the bulge stellar population to be pressure supported with a sub-dominant cylindrical rotational component to provide the pattern speed of the bar. In the model, however, the quasi-circular velocity component, together with the range of bulge star Galactocentric distances, leads to this velocity component dominating the longitudinal proper motion dispersion. The dichotomy in near-side bulge versus far-side bulge velocities imposed by a large circular velocity also increases the event rate of events with far-bulge sources and near-bulge lenses, which have the largest Einstein radii, so will increase the mean Einstein radius of bulge–bulge lenses somewhat. The increased longitudinal proper motion dispersion of bulge stars will also increase the relative event rate of bulge–disk lensing. In all cases, BGM1106's event timescales will be shorter than would be produced by more realistic kinematics.

6.2.3. Bulge Initial Mass Function

As discussed already in Section 3.2.1, BGM1106's mass function in the bulge differs from typically assumed mass functions (e.g., Kroupa 2001) in both its shallow slope $dN/dM \propto M^{-1.0}$ and its high lower-mass cutoff of $0.15 M_\odot$. We found that replacing the mass function with something more reasonable, such as from Sumi et al. (2011), a (Kroupa 2001) slope of $M^{-1.3}$ between 0.08 and $0.7 M_\odot$ and a slope of $M^{-0.5}$ between 0.01 and $0.08 M_\odot$, would improve the match of the shape of BGM1106's luminosity function to measurements from Calamida et al. (2015). Wegg et al. (2017) find that a similar mass function, when combined with an N -body bulge model fit to infrared star counts and RV distributions, can simultaneously fit both microlensing optical depths and timescale distributions, though the slope of the brown-dwarf mass function in the bulge remains uncertain (-0.65 ± 0.89). We note that we have not considered the effects of age or metallicity that can affect star counts, especially for evolved stars, or in the case of metallicity, M dwarfs as well.

Adding stars and brown dwarfs below $0.15 M_\odot$ to the BGM mass function would increase the optical depth and event rate per star for bulge–bulge lensing by factors of ~ 1.9 and ~ 3.4 , respectively. The disk's mass function in BGM1106 has a more typical slope, and extends down to $0.08 M_\odot$ but adding brown dwarfs would increase the optical depth and event rate for bulge–disk lensing also, but to a lesser degree than for bulge–bulge lensing. It is therefore likely that the form of the mass function can explain a significant amount of the factor of ~ 2.1 underprediction of the event rate by BGM1106, the factor of ~ 1.8 underprediction of the optical depth, and its shallower slope of the luminosity function. Adding low-mass stars to the mass function would also act to decrease the mean event timescale (Awiphan et al. 2016).

6.2.4. Extinction

The BGM1106 uses the 3D extinction model of Marshall et al. (2006) to provide extinction as a function of distance, and the Cardelli et al. (1989) reddening law with $R_V = 3.1$ to convert extinctions in the K_s -band to other wavelengths.

Schultheis et al. (2014) assessed the performance of various 2D and 3D extinction maps and found that generally 3D extinction maps were accurate, but failed along certain sight lines. Numerous studies have shown that the reddening law toward the bulge differs from the $R_V = 3.1$, with $R_V \sim 2.5$ more typical (e.g., Nataf et al. 2013). It is even possible that the reddening law deviates from a power law in the $1\text{--}2 \mu\text{m}$ range (Hosek et al. 2018).

Despite this uncertainty, it is likely that the impact of errors in the extinction will only have a small impact on the predicted yields. This is simply because the total extinction across our fields is only $A_H \sim 0.6$, though it does reach $A_H \sim 1$ in small parts of the fields closest to the plane. Therefore, even a large fractional error corresponds to a relatively small absolute error. We can use the analytic framework in Section 5.1 to estimate the impact. Errors in the extinction will affect all stellar noise (source and blends) equally. A 33% underestimate of the extinction, or 0.2 mag , reduces flux and $\Delta\chi^2$ by 17%. Using the slope of the $\Delta\chi^2$ distribution from Table 6, $\alpha = -0.399$ for $1 M_\oplus$ planets, and the Cycle 7 design, the underprediction of extinction would result in an overprediction of the Earth-mass planet yield of $\sim 7\%$. Averaged over the whole proposed fields, an error this large seems unlikely. However, we note that the relatively coarse extinction map we have used (resolution $0^\circ.25 \times 0^\circ.25$) may have impacted our field optimization.

6.2.5. Impact on WFIRST's Planet Yield

In Section 3.2 we have adopted an event rate scaling to match measured microlensing event rates per red clump source and the number of faint sources. This correction factor will be valid to first order because the majority of WFIRST's sources will be in the bulge, like the red clump stars. Therefore, the issues raised with the model in this section should not affect our predictions for the number of microlensing events WFIRST will detect. However, they will affect the properties of the events, which may impact the detection efficiency.

The lack of low-mass stars and brown dwarfs means the mean lens mass is too large, and at fixed planet mass the mass ratio will be too small. This has the effect of reducing our detection efficiency per event slightly at fixed planet mass. BGM1106's bulge kinematics result in timescales that are too short, which results in reduced planet detection efficiency due to shorter planetary anomalies. The increase in Einstein ring radii due to the elongated bulge will increase timescales slightly, but not enough to counteract the effect from kinematics.

In addition to the average detection efficiency, the mass function and bar angle issues cause BGM1106 to overestimate the average Einstein radius. This means that its peak sensitivity to planets will be at a slightly smaller semimajor axis than indicated by Figure 9.

We have concluded that BGM1106 probably underpredicts the number of source stars because of its shallow mass function slope. However, we corrected for the underprediction by simply multiplying the event rate for all source stars, which likely has the effect of overcorrecting for bright stars, brighter than $F814W_{\text{Vega}} \sim 22$. Roughly two thirds of $1 M_\oplus$ planet detections come from events with source stars above this boundary, which would imply that the an overcorrection of bright source stars leads to an over estimate of $\sim 14\%$ in Earth-mass planet yield, and probably a larger overestimate for the most difficult to detect planets. However, if we had estimated

our event rate scaling using the MOA all-star event rates instead of the ERC event rates (see Figure 6) and only used data in the range of Galactic latitudes where *WFIRST* will probably observe, we would have derived a larger correction.

In summary, the issues with the Galactic model after corrections act to both increase and reduce the detection efficiency or number of events by mostly small factors. To a certain degree, then, we can expect the effects of different signs to cancel, and the associated uncertainties to grow in quadrature. It is likely therefore that a single large uncertainty will dominate over smaller uncertainties. The quantity with the largest uncertainty is probably the microlensing event rate, due to the relatively low S/N when subdivided down to square-degree scales, and the need to extrapolate closer to the Galactic plane. We reiterate, however, that while there will remain significant uncertainties in the absolute yield predictions, the relative yields between designs simulated with the same methodology and common parameters and assumptions will be much less uncertain.

6.3. Planet Population Uncertainties

The uncertainty in our assumptions about the population of planets is large, and some regions of the parameter space are completely unconstrained. A major goal of the survey, after all, is to detect and measure the occurrence rate of the cold planet population that cannot be conducted by any other method, or by microlensing observations from Earth. Nevertheless there are measurements of planet abundances in the regions of *WFIRST*'s sensitivity from microlensing surveys (Gould et al. 2010; Sumi et al. 2010; Cassan et al. 2012; Clanton & Gaudi 2016; Shvartzvald et al. 2016; Suzuki et al. 2016), and near it from transits (e.g., Burke et al. 2015), RV (e.g., Cumming et al. 2008; Johnson et al. 2010; Bonfils et al. 2013; Montet et al. 2014), and direct imaging (Nielsen & Close 2010; Bowler et al. 2015; Chauvin et al. 2015). These observational constraints allow us to anchor extrapolations into the regions that are unexplored.

Our fiducial joint mass–semimajor axis occurrence distribution (in many places we have simply referred to this as our fiducial mass function) assumes a broken power law in mass, and a log-uniform distribution in semimajor axis. The high-mass end of the mass function was chosen to match the estimate of the mass function from Cassan et al. (2012) based on microlensing searches. This measured mass function only extends down to $\sim 5 M_{\oplus}$, so we chose to saturate the power law at a value of two planets per dex mass per dex semimajor axis per star at $5.2 M_{\oplus}$ to prevent overly optimistic predictions of large numbers of low-mass planets. Since adopting this as our fiducial mass function, several studies have advanced upon the Cassan et al. (2012) result on which it was based. While we have chosen to retain the fiducial mass function for consistency and easy comparison with past *WFIRST* reports, it is worth examining what impact that adopting another joint mass–semimajor axis function would have.

Suzuki et al. (2016) analyzed a larger and more homogeneously selected (compared to Cassan et al. 2012) set of microlensing planet detections from the MOA survey. Shvartzvald et al. (2016) also studied the mass ratio function with a smaller sample of events. Both mass ratio function measurements are shown in Figure 21 in comparison to our fiducial mass function. To convert mass ratio to planet mass we assumed a host mass of $0.5 M_{\odot}$. Suzuki et al. (2016) found

Figure 21. Comparison of our fiducial mass function (solid gray line) to the latest measurements based on microlensing data. Brown and green points with error bars show the planet occurrence as a function of mass ratio converted into planet masses, assuming a $0.5 M_{\odot}$ host star, from Suzuki et al. (2016) and Shvartzvald et al. (2016), respectively, and the dotted gray line shows a fit to microlensing, radial velocity trends, and direct imaging data by Clanton & Gaudi (2016). Dark blue data points show the mass function measurement precision of a mock survey by *WFIRST* (Cycle 7 design) in 0.5 dex mass bins assuming only half of the planet detections can be used; filled points follow the fiducial mass function, while open points follow a mass function that saturates at an occurrence rate a factor of 4 lower than the fiducial mass function.

evidence for a turnover in the planet mass *ratio* function at $q \approx 1.7 \times 10^{-4}$, which corresponds to a planet mass of $\sim 30 M_{\oplus}$ assuming a mean host mass of $0.5 M_{\odot}$. This turnover appears to be confirmed by an independent analysis using very different methods (Udalski et al. 2018), but the exact shape of the turnover, and the steepness of the subsequent decline below this mass, are very poorly constrained. In a bin centered at $M_p \approx 3 M_{\oplus}$ Suzuki et al. (2016) place a 95% upper limit on the mass function of 0.5 planets per dex², about a factor of four below the value of our saturated fiducial mass function. At ~ 30 – $100 M_{\oplus}$ the fiducial mass function is a good match to the Suzuki et al. data, but at larger masses it again overestimates the measurements, but to a lesser degree than at small masses. Altogether, this suggests that our total yields will be overestimated somewhat, and the yields at low masses could be overestimated significantly.

It is important to recognize that the planet yields assuming a fiducial mass function that we estimate in this study are only a useful proxy for the true value of the survey, which is planet sensitivity, or its power to measure the planet occurrence rate. This sensitivity is independent of the assumed mass function and, as can be seen from Figure 9, it will extend down to $\sim \text{few} \times 10^{-2} M_{\oplus}$. In this sense the sensitivity range is the point at which no detection of planets of mass M_p during the mission ceases to be an interesting constraint on planet occurrence. In Figure 21 we show that, even with a mass function saturation value a factor of 4 lower than our fiducial mass function, *WFIRST* will detect a sufficient number of planets to measure the mass function in 0.5 dex bins to below $1 M_{\oplus}$, and it would set interesting upper limits on the planet occurrence at masses below this. In these estimates we assumed that only half of *WFIRST*'s planets would be utilized in this mass function measurement. Over its entire range it will provide measurements of the mass function with far greater precision than is currently possible from ground-based surveys. While not a direct comparison, Henderson et al. (2014a) predict that

KMTNet, during its nominal five year survey, would find a factor of ~ 16 fewer $1 M_{\oplus}$ planets than *WFIRST* will find in its nominal (Cycle 7) survey.

In addition to uncertainties in the planet mass function, there is even greater uncertainty in the form of the planet occurrence as a function of semimajor axis near to *WFIRST*'s peak sensitivity. Clanton & Gaudi (2014a) have shown that there is at present very little overlap in the sensitivity regions of current microlensing and RV surveys, but RV surveys tend to show an increase in planet occurrence with log semimajor axis or log period (e.g., Cumming et al. 2008; Bonfils et al. 2013) that appears to be consistent with microlensing occurrence rates when extrapolated (see, e.g., Gould et al. 2010; Suzuki et al. 2016). Results from *Kepler* show a similar rising trend for large planets, but a shallow decline in occurrence beyond $P \sim 10$ days for planets smaller than Neptune (e.g., Petigura et al. 2018); whether these trends continue throughout *WFIRST*'s region of sensitivity is unconstrained at present. At larger orbital separations, Clanton & Gaudi (2016) find that a cut-off in planet occurrence at ~ 20 au is required to remain consistent with microlensing, RV trends, and direct imaging results.

6.4. Future Improvements

To make further progress on estimating the yields of the *WFIRST* survey requires work on several fronts. The most critical need is observational, due to the long lead time necessary to observe, analyze, and interpret new data. Advances in simulations are also needed to better understand the relative importance of mass measurements in optimizing *WFIRST*'s fields and observing strategy.

Observationally, the most important measurement to make is of the microlensing event rate in the potential *WFIRST* fields, in the infrared and to a depth as close as possible to that achievable by *WFIRST*. Such a measurement will also test the ability of event rate models based only on star counts (Poleski et al. 2016) to predict event rates closer to the Galactic plane. This requires an infrared microlensing survey in order to penetrate the dust near the Galactic plane and reach sources throughout the bulge and into the far side of the disk. Such a survey is underway using the UK Infrared Telescope (UKIRT, Shvartzvald et al. 2017, 2018), but VISTA, an infrared telescope with a better location for bulge observations and a field of view that is around three times larger than UKIRT's, is not currently conducting observations optimized for microlensing (though note that the VVV survey has discovered a number of microlensing events; Navarro et al. 2017). In the time it takes the UKIRT survey to build up enough events, progress can be made by analyzing the full data sets of MOA and OGLE surveys. Currently the study measuring optical depths and event rates with the largest number of events analyzes 474 MOA events from the 2007 and 2008 seasons (Sumi et al. 2013; Sumi & Penny 2016; Wyrzykowski et al. 2015 only provides the timescale distribution of a larger sample of 3718 events from the OGLE-III survey, and not optical depth and event rate measurements). The current phase of the MOA survey has now been operating for over a decade, and the OGLE-IV survey has been discovering ~ 2000 events a year since 2011. Analysis of the full data sets of both these surveys would enable measurements of the event rate at much higher spatial resolution than is now possible.

In addition to the event rate, a better understanding is needed of the source magnitude distribution in the infrared. Deep

luminosity functions in the *I* band are available at latitudes of $b \approx -6, -3.8$ and -2.7 (Holtzman et al. 1998; Zoccali et al. 1999; Calamida et al. 2015), and in the *J* band at $b \approx -6$ (Zoccali et al. 1999). *WFIRST* will probably observe much closer to the plane and, in addition to bulge stars, there will be a significant contribution from stars in the near and far disk, which will have very different event rates. Understanding the breakdown of components will require new, deep, infrared magnitude distributions and proper motion measurements from *Hubble* and/or the *James Webb Space Telescope* (*JWST*), ideally for several sight lines in the potential *WFIRST* fields. Images produced from special high-density mode scans with *Gaia*¹⁰ can cover a much larger area than *HST* or *JWST*, but so far have only been carried out in Baade's window, too far from the expected *WFIRST* fields. Every high-resolution image taken in the *WFIRST* fields before launch will also provide a "precovery" data set for lens mass measurements, extending the baseline over which PSF elongation and color-dependent centroid shifts (e.g., Bennett et al. 2007) can be measured. A small sample of these could be used as ground truth for the *WFIRST* measurements over the 4.5 yr mission baseline. An ambitious survey of the entire $\sim 2 \text{ deg}^2$ *WFIRST* microlensing survey fields, in a similar manner to the Panchromatic Hubble Andromeda Treasury survey (Dalcanton et al. 2012), would require ~ 1500 *HST* pointings, or ~ 750 *JWST* NIRCAM pointings, and would provide immense legacy value for what will become one of the most intensely observed patches of sky (Yee et al. 2014). For example, optical photometry from the *HST* could provide photometric metallicity estimates for every source and a large fraction of lens stars. *JWST* 3.6 and $4.5 \mu\text{m}$ imaging could be used to measure star-by-star extinctions using the Rayleigh-Jeans color excess (e.g., Majewski et al. 2011). Both the *JWST* and *HST* will have roughly twice the angular resolution of *WFIRST* and so their imaging can assist in cases where local, random stellar overdensities hamper the interpretation of *WFIRST* images. The improved resolution and increased time baseline would vastly improve measurements of proper motions for all stars in the *WFIRST* field, which would help in the measurement of parallaxes from *WFIRST*'s microlensing data (Gould et al. 2014a).

There remains significant room for improvement in Galactic models. Fully simulating a *WFIRST*-like survey requires all the features of a population synthesis model, in order to understand not just the event rates, but the properties of the lenses and sources. The Besançon model is presently the only publicly accessible population synthesis model that incorporates kinematics. The latest version of the model¹¹ has some changes relative to the version we used here, but many of the problems we have identified with it remain (e.g., the small bar angle, too-fast kinematics). The publicly accessible Besançon model also lacks any flexibility to adjust model parameters, which is important for maintaining a model in agreement with burgeoning data sets and for understanding the propagation of model uncertainties to yields. The TRILEGAL model (Girardi et al. 2005; Vanhollebeke et al. 2009) provides some flexibility to adjust structural parameters, but the publicly available version does not include kinematics.¹² The publicly available Galaxia

¹⁰ https://www.esa.int/spaceinimages/Images/2017/08/Gaia_sky_mapper_image_near_the_Galactic_centre

¹¹ <http://modele2016.obs-besancon.fr/>

¹² <http://stev.oapd.inaf.it/cgi-bin/trilegal>

code (Sharma et al. 2011) implements a version of the Besançon model and can also accommodate N -body models, potentially providing the necessary flexibility. New, more flexible versions of the Besançon model are under development, that improve the evolutionary tracks, add flexibility of the IMF, star formation history, and bar angle (e.g., Czekaj et al. 2014; Lagarde et al. 2017). GALMOD is another new population synthesis model accessed via web forms with significant flexibility in its model parameters (Pasetto et al. 2018).

There is also significant work needed on microlensing simulations to better understand the information that *WFIRST* will be able to measure for each planet it finds. This is especially the case for host mass measurements, which will be possible though one or more of the techniques: detecting the host as it separates from the source and measuring image elongation, color-dependent centroid shifts or directly resolving the lens (e.g., Bennett et al. 2007; Henderson 2015; Bhattacharya et al. 2017), measuring the microlensing parallax with or without finite source measurements (e.g., Yee et al. 2013; Yee 2015; Bachelet et al. 2018), or even measuring astrometric microlensing (Gould & Yee 2014). The error budget of these measurements is likely to be dominated by systematic errors, and so more detailed end-to-end simulations of the stacking, photometry, and astrometry pipelines are likely necessary in order to fully understand *WFIRST*'s capabilities. These simulations will then allow the optimization of *WFIRST*'s survey for characterized planets and not just the total number of detected planets. It will also be important to understand how *WFIRST* performs for more exotic planetary systems (e.g., multiplanet systems, Zhu et al. 2014; circumbinary planets, Luhn et al. 2016; exomoons, Liebig & Wambsganss 2010; etc.) and for rejecting possible false-positive detections caused by, for example, binary source stars (e.g., Gaudi 1998).

7. Conclusion

We have performed detailed simulations of several potential designs of the *WFIRST* mission in order to estimate the planet detection yield of its microlensing survey. We derived a correction factor to apply to microlensing event rates computed using the Besançon model in order to normalize event rates to those measured by microlensing surveys. Having done so, we estimate that the most recent *WFIRST* design (Cycle 7) will be able to detect ~ 180 Earth-mass planets and ~ 1400 cold exoplanets in total. For Earth-mass planets, its sensitivity will extend from ~ 1 au outwards, and will have a wider range of sensitivity at higher masses. The lower limit of *WFIRST*'s sensitivity for planets in suitable orbits extends down to the mass regime of the solar system's moons, e.g., Ganymede. The mission will fulfil the goals assigned to its microlensing component by the 2010 decadal survey committee to "determine how common Earth-like planets are over a wide range of orbital parameters." However, significant observational, Galactic modeling, and simulation work still needs to be done in order to optimize and fully understand the yields of the survey.

This paper is dedicated to the memory of Neil Gehrels who, as *WFIRST* Project Scientist, gracefully shepherded this mission from its formal initiation in 2016, until his untimely death. We miss him and his leadership dearly. We would also like to thank members of the *WFIRST* Microlensing Science

Investigation Team, as well as Jeff Kruk, Dave Content, Kevin Grady, and many others in the *WFIRST* program office for their support. We would like to thank Chris Hirata and Jay Anderson for their practical help and discussions. We thank the referee for their efforts. Finally, we would like to single out David Bennett, who has been a constant source of input and guidance as we have developed our models and worked on this paper. We recognize that, without his tireless efforts, a microlensing survey on *WFIRST* would not exist.

This work was performed in part under contract with the California Institute of Technology (Caltech)/Jet Propulsion Laboratory (JPL) funded by NASA through the Sagan Fellowship Program executed by the NASA Exoplanet Science Institute. M.T.P. and B.S.G. were supported by NASA grants NNX14AF63G and NNG16PJ32C, as well as the Thomas Jefferson Chair for Discovery and Space Exploration. S.M. was supported by NSFC grants No. 11333003, 11390372, and 11761131004.

Facility: Exoplanet Archive.

Software: Matplotlib (Hunter 2007), NumPy (Oliphant 2006), SciPy (Jones et al. 2001), Astropy (The Astropy Collaboration et al. 2013), gnuplot, WebbPSF (Perrin et al. 2012), General Polygon Clipper library (Murta 2015), MATLAB package for astronomy and astrophysics (Ofek 2014).

Appendix A Improvements to GULLS

A.1. Custom Bandpasses

WFIRST will perform its microlensing survey using a wide infrared filter, covering the full range of detector sensitivity in order to maximize photon count rates. The Besançon model provides stellar magnitudes in several standard photometric systems, computed from stellar atmosphere models (Robin et al. 2003), but it is not easy to compute magnitudes in additional bands. We solved the problem of calculating magnitudes in custom bandpasses by producing what amount to smoothed spectral energy distributions (SEDs) by interpolating between the different available pass bands, and then integrating the product of the smoothed SED with the system throughput curve for each desired custom filter.

The BGM only outputs a limited number of stellar magnitudes out of the whole range available—in the catalogs we were using R , I , J and H —so we found it necessary to supplement these with synthesized K and L magnitudes in order to completely cover and bracket the *WFIRST* detector sensitivity range. The magnitudes in these bands were synthesized by assuming that the star was emitting as a blackbody, and that extinction followed the extinction law listed in the BGM1106 header data, namely $A_K/A_V = 0.118$ and $A_L/A_V = 0.0$. The smoothed SEDs were interpolated using radial basis functions (RBFs).

To test the synthetic *WFIRST* magnitudes we combined catalogs from three BGM sight lines ($\ell = -0.4^\circ$ and $b = -3.2^\circ, -1.95^\circ$, and -0.7°) and mean extinctions to bulge stars of $A_V = 2.35, 3.62$, and 14.37 , respectively. Colors and magnitudes were dereddened using values of $A_{W149}/A_V = 0.225$ and $E(Z087 - W149)/A_V = 0.208$, chosen to align the red clump of each field. These values compare to 0.210 and 0.295, respectively, computed at the central wavelength of the filters using the Cardelli et al. (1989) extinction law with

Figure 22. Color-absolute magnitude diagram in the principal *WFIRST* microlensing filters, W149 and Z087. Shown are stars from three BGM1106 sightlines at $\ell = -0^\circ.4$ and $b = -3^\circ.2$, $-1^\circ.95$, and $-0^\circ.7$, combined, with disk stars plotted with black dots and bulge stars with red dots. The evolutionary tracks and synthetic colors differ between the disk and bulge stars as described by Robin et al. (2003) and references therein. The blue line shows a MIST version 0.3 (Choi et al. 2016) isochrone for a 10 Gyr, $[\text{Fe}/\text{H}] = 0.0$ population computed using *WFIRST* filter profiles, shifted by $\Delta(Z087 - W149) = -0.05$, demonstrating that our scheme for computing *WFIRST* colors and magnitudes works reasonably well. To aid conversion to apparent magnitudes, the distance modulus to the bulge population is approximately 14.5, and the extinction and reddening will typically be $A_{W149} \sim 0.5$ and $E(Z087 - W149) \sim 0.5$ in the expected *WFIRST* microlensing fields.

$R_V = 3.1$; we can expect some difference due to the very wide bandpass of the W149 filter. Figure 22 shows the resulting color-absolute magnitude diagram. The dereddening does not work perfectly, with the turnoff location of bulge stars differing by ~ 0.03 mag in $Z087 - W149$ for the different values of extinction. We note that disk stars tend to be redder than the bulge stars, likely due to the different stellar evolution and synthetic photometry models used for each population as described by Robin et al. (2003) and references therein. We also compared the dereddened BGM stars to a MIST version 0.3 isochrone computed using *WFIRST* bandpasses (Choi et al. 2016) for a 10 Gyr, solar metallicity population. Subtracting 0.05 mag from $Z087 - W149$ MIST isochrone brings it into good alignment with the main-sequence BGM1106 stars, though the colors of the turn-off and giant branch disagree by ~ 0.05 mag in opposite directions; it is possible that differences in stellar evolution codes or the filter transmission curves used for the MIST isochrones could cause these problems. Overall, it appears that our scheme for computing magnitudes in the *WFIRST* bandpasses is reasonable, and should not inject errors significantly larger than the theoretical uncertainties associated with the choice of isochrones.

A.2. Zodiacal Light Model

Outside the Earth’s atmosphere, in the near-infrared, the brightest diffuse background is caused by the zodiacal light: light from the Sun scattered off interplanetary dust grains. In P13 we assumed this was constant, taking the mean value at the times that *Euclid* could possibly conduct a putative bulge microlensing survey (see P13 for details of this survey). However, the level of the zodiacal background varies as a function of the elongation of the target fields relative to the

Sun, an effect that may become important for long observing seasons. For this reason we incorporate a full-sky model of the zodiacal light, removing the need for the GULLS user to calculate the average level of the zodiacal light in their required bandpasses.

The zodiacal light brightness in a given bandpass is calculated by integrating the RBF interpolated zodiacal light spectrum (as provided by Leinert et al. 1998, including solar elongation-dependent color terms) over the throughput curve of the bandpass. The spatial dependence of the zodiacal light is calculated by RBF interpolation of the map provided by Leinert et al. (1998).

A.3. Faster Photometry Routines

In P13 we performed photometry on a pixel-by-pixel noise realization of each image at each epoch. This was computationally expensive, and in certain circumstances was the primary bottleneck of the computation. To speed up the photometry we implemented a routine that takes as input a noiseless realization of the baseline image, accounts for blending, and returns a simple function to compute the photometric signal and uncertainty as a function of magnification. With this, only a single realization of the noise on the photometric data point is needed. The routine also solves for the magnification at which saturation is reached in one of the pixels of the aperture, allowing saturation to be identified accurately without building a realization of an image.

A.4. Improved Observer-centric Velocities/Timescales

In order to include parallax effects in light curves we now compute geocentric or, more accurately, observer-centric, microlensing event timescales. In (P13), we used heliocentric velocities to compute event timescales. Due to the observer’s motion about the Sun (typically the Earth’s, which ranges from -30 to 30 km s^{-1} projected onto the sky), the relative source-lens velocity will change over the course of the year. This compares to the ~ 200 – 1000 km s^{-1} projected velocities of the source and lens (see, e.g., Calchi Novati et al. 2015), implying a typical modification of the timescales by $\sim 3\%$ – 15% . However, for both *Euclid* and *WFIRST*, microlensing observations will be made at or near quadrature, meaning that the projected velocity of the Earth will be close to zero. The effect on planet yields of using the improved timescales is therefore likely to be very small. Some of the simulations presented here use the improved timescales, while others were completed before the improved timescale calculation was implemented. Full details are given by Penny et al. (2017).

Appendix B

Computing the Mass-Semimajor Axis Sensitivity Curve or “The Making of Figure 9”

In order to compute the sensitivity curve shown in Figure 9 required computing the planet detection rate on a grid of planet mass and semimajor axis, spaced by 0.25 and 0.125 dex, respectively. Obtaining reasonably accurate results is computationally intensive, with the required computation increasing as one over the square of the detection efficiency in order to achieve equal Poisson statistical uncertainties. With t_0 drawn from any point in the five year mission, and a $\sim 24\%$ observing duty cycle, the detection efficiency is $\sim 10^{-5}$ at the three-detection line. This implies that a 10% statistical error would

Figure 23. Examples of light curves exhibiting numerical errors that were not caught by error reporting routines. Each type of error is discussed in the text. Data points with error bars show the predicted magnitude and uncertainty of measurements without any noise, the blue line shows the simulated event, including numerical errors, and the green line shows the best-fit single lens model. Parameters for each event shown are listed above each plot; α is the angle subtended by the source trajectory relative to the binary axis.

require $\sim 100/10^{-5} = 10^7$ light curves to be generated at each grid point near the sensitivity curve, with most of these showing no detection. To make the computation tractable we developed the CROIN parameterization (Penny 2014), which is used to generate only light curves where there is a reasonable probability of a planet detection. This coordinate system is centered on the planetary caustic(s), and the region around the caustic that contains a detectable planetary signature is a circle of radius $r_c(s, q)$, whose analytic functional dependence on the projected separation relative to the Einstein radius s and mass ratio q was derived empirically by Penny (2014). Only source trajectories with impact parameters relative to the planetary caustic $u_c < r_c$ are simulated. We used the CROIN parameterization for planet masses $M \leq 10 M_\oplus$ and for larger masses with $\log(a/\text{au}) \geq 1.125$, reducing the number of required light-curve computations by more than two orders of magnitude. When using the CROIN parameterization we still require that the main-event impact parameter and peak time obey $-3 \leq u_0 < 3$ and $0 \leq t_0 < 2011$.

For the low-mass planets that *WFIRST* is sensitive to, computing the light curve is not a trivial operation and is prone to numerical errors. For its speed, we primarily relied on a contour integration code (Gould & Gaucherel 1997; Dominik 1998)

written by S. Mao. This solves the complex fifth-order binary lens polynomial at many points around the source circumference and then links the resulting solutions into a number of potentially merging images. The coefficients of the polynomial have additive terms of the order of 1, and various combinations of powers of q and s . For the lowest-mass planets we consider, $q \sim 10^{-8}$, so we suspect that the numerical errors are a result of catastrophic cancellations in parts of the calculation where this is difficult, if not impossible, to avoid. The vast majority of errors are caught by error handling routines, and when this occurs the light curve is passed to a much slower but more robust inverse ray shooting routine (e.g., Kayser et al. 1986), but occasionally errors slip past the error handling routines.

In order to make sure that these numerical errors were not significantly affecting results we visually inspected a large sample of the light curves of the lowest-mass exoplanet detections. We found examples of errors that *did* cause false-positive detections and *could* cause false-negative detections. As our simulations only output the light curves of a sample of planet *detections*, we are not able to assess the degree to which our predicted yields are reduced by false negatives. As we correct for false positives (see below) and not for false negatives our planet yield predictions at the lowest masses are

Figure 24. Fraction of false-positive planet detections for points on the planet mass–semimajor axis grid. In each cell the fraction is the number of false-positive light curves over the number of inspected light curves. The number below that is the percentage of false positives, weighted by the event rate as described in Section 3. The shading is a linear scale that saturates to yellow at three detections during the survey, before correction for the false-positive rate.

likely to be conservative. Figure 23 shows four examples that represent the overwhelming majority of the errors we found—two were false positives, and two were errors that did not affect the designation of the event as a detection, but had the errors been more severe one of these would have resulted in a false negative.

The first example (a) shows a type of false positive that only occurs in high-magnification events where the source is resolved by the magnification structure surrounding the host star. Most of these examples had only a single discrepant data point. However, due to the high photometric precision that high-magnification events enable, the discrepant point causes a large $\Delta\chi^2$.

The second example of false positive (b) is one of a more general class where a significantly discrepant data point (or several) occurs during a planetary anomaly. They can be either positive or negative deviations, and are typically sharp changes relative to the source crossing time. The events with these numerical errors are only classified as false positives if removal of the discrepant points would move the event below the $\Delta\chi^2$ threshold.

A potential false-negative non-detection (c), which we call a “finite-source drop-out,” occurs for events with wide-separation planets that show a small amplitude top-hat planetary signature due to extreme finite-source effects. They are caused when two images near the planet are both incorrectly flagged as false solutions to the lens equation, leaving three valid solutions instead of five.¹³ In principle it is possible for all the data points during a planetary anomaly to experience this problem, in which case the event would not count as a planet detection and the simulation would not output the light curve. It is therefore

¹³ The binary lens equation will have either three or five solutions depending on whether the source is inside or outside a caustic, but the complex fifth-order polynomial that can be formed by rearranging the lens equation always has five solutions. The validity of each solution is checked by inverse shooting a ray from each candidate solution and checking the proximity of the ray to the source position from which it should have originated.

Figure 25. Average rate-weighted false-positive fraction as a function of planet mass. Filled black points show the rate-weighted false-positive fraction as measured by inspecting light curves. The open red points show the same data but showing the quantity $\log(1/\text{FPF} - 1)$ plotted against the axis on the right. The lines show the same best-fit analytic model.

impossible to estimate the number of these occurrences by our current method of inspection. However, we can guess that the number is likely small because most instances of drop-outs show only a small fraction of the top hat dropping out.

The final example (d) of numerical errors is likely very similar to the finite source drop-out but occurs for all separations. Again, the magnification is artificially reduced, but usually by an insignificant amount that does not affect the event’s status as a detection or non-detection.

The false-positive fraction becomes significant as the planet mass decreases, so it is important that we correct for it. Figure 24 shows the event-rate-weighted (see Section 3)

Figure 26. False-positive-corrected planet detection rates plotted as a function of mass for each of the semimajor axis grid points. The detection rates have been shifted vertically by an amount $10 \log(a/\text{au})$, and are also color-coded by semimajor axis with black corresponding to $a = 0.01$ au and yellow to $a = 100$ au. Curved lines are quadratic fits to the points with detection rates within a factor of 30 of the three detections per survey threshold that is plotted in Figure 9. For each semimajor axis we plot the three-detection threshold as a horizontal line and an open square point marking the position at which the quadratic curve crosses the threshold.

false-positive fraction for points on the planet mass–semimajor axis grid, along with the number of events inspected and the number of events that were false positives. The initial intention was to correct each point on the grid individually by its own false-positive rate, but there were not enough light curves output by the simulation for this to be accurate. Instead, given that the false-positive rate seemed to be relatively independent of semimajor axis (within the large error bars), we took the rate-weighted average of all semimajor axes at each value of the mass. The false-positive fractions can be transformed to a form that is roughly a single power law in mass

$$\log\left(\frac{1}{\text{FPF}} - 1\right) = \alpha \log(M/M_{\oplus}) + \beta, \quad (16)$$

where FPF is the false-positive fraction and $\alpha = 1.25$ and $\beta = 2.7$ are the best-fit linear regression parameters. This is shown in Figure 25.

Returning to the problem of estimating *WFIRST*'s sensitivity in the mass–semimajor axis plane, we correct the gridded planet detection rates using Equation (16). For a given semimajor axis, to find the mass at which the planet detection rate crosses the three-detection threshold we fit a quadratic polynomial in $\log(M/M_{\oplus})$ to the log of the planet detection rates that are within a factor of 30 of the three-per-survey threshold, and then solve the resulting quadratic equation (taking care to select the appropriate root). We repeat this process for each semimajor axis grid point as is shown in Figure 26. We excluded the points with $\log(a/\text{au}) = -2$ and -1.875 from the analysis because our grid did not extend to high enough masses to properly bracket the point at which three detections are expected.

ORCID iDs

Matthew T. Penny <https://orcid.org/0000-0001-7506-5640>
 B. Scott Gaudi <https://orcid.org/0000-0003-0395-9869>
 Eamonn Kerins <https://orcid.org/0000-0002-1743-4468>
 Nicholas J. Rattenbury <https://orcid.org/0000-0001-5069-319X>
 Shude Mao <https://orcid.org/0000-0001-8317-2788>
 Sebastiano Calchi Novati <https://orcid.org/0000-0002-7669-1069>

References

- Akeson, R. L., Chen, X., Ciardi, D., et al. 2013, *PASP*, **125**, 989
 Awiphan, S., Kerins, E., & Robin, A. C. 2016, *MNRAS*, **456**, 1666
 Bachelet, E., Hinse, T. C., & Street, R. 2018, *AJ*, **155**, 191
 Barclay, T., Quintana, E. V., Raymond, S. N., & Penny, M. T. 2017, *ApJ*, **841**, 86
 Beaulieu, J. P., Kerins, E., Mao, S., et al. 2008, arXiv:0808.0005
 Bennett, D. P., Anderson, J., Beaulieu, J. P., et al. 2010a, arXiv:1012.4486
 Bennett, D. P., Anderson, J., & Gaudi, B. S. 2007, *ApJ*, **660**, 781
 Bennett, D. P., Bally, J., Bond, I., et al. 2003, *Proc. SPIE*, **4854**, 141
 Bennett, D. P., Bhattacharya, A., Anderson, J., et al. 2015, *ApJ*, **808**, 169
 Bennett, D. P., & Rhie, S. H. 1996, *ApJ*, **472**, 660
 Bennett, D. P., & Rhie, S. H. 2002, *ApJ*, **574**, 985
 Bennett, D. P., Rhie, S. H., Nikolaev, S., et al. 2010b, *ApJ*, **713**, 837
 Bennett, D. P., Rhie, S. H., Udalski, A., et al. 2016, *AJ*, **152**, 125
 Bhattacharya, A., Beaulieu, J. P., Bennett, D. P., et al. 2018, *AJ*, **156**, 289
 Bhattacharya, A., Bennett, D. P., Anderson, J., et al. 2017, *AJ*, **154**, 59
 Bienayme, O., Robin, A. C., & Creze, M. 1987, *A&A*, **180**, 94
 Bolcar, M. R. 2017, *Proc. SPIE*, **10398**, 10
 Bond, I. A., Bennett, D. P., Sumi, T., et al. 2017, *MNRAS*, **469**, 2434
 Bonfils, X., Delfosse, X., Udry, S., et al. 2013, *A&A*, **549**, 109
 Borucki, W. J., Koch, D. G., Lissauer, J. J., et al. 2003, *Proc. SPIE*, **4854**, 129
 Bowler, B. P. 2016, *PASP*, **128**, 102001
 Bowler, B. P., Shkolnik, E. L., & Tamura, M. 2015, *ApJS*, **216**, 7
 Burke, C. J., Christiansen, J. L., Mullally, F., et al. 2015, *ApJ*, **809**, 8
 Butler, R. P., Vogt, S. S., Marcy, G. W., et al. 2004, *ApJ*, **617**, 580
 Calamida, A., Sahu, K. C., Casertano, S., et al. 2015, *ApJ*, **810**, 8
 Calchi Novati, S., Gould, A., Udalski, A., et al. 2015, *ApJ*, **804**, 20
 Campbell, B., Walker, G. A. H., & Yang, S. 1988, *ApJ*, **331**, 902
 Cao, L., Mao, S., Nataf, D., Rattenbury, N. J., & Gould, A. 2013, *MNRAS*, **434**, 595
 Cardelli, J. A., Clayton, G. C., & Mathis, J. S. 1989, *ApJ*, **345**, 245
 Cassan, A., Kubas, D., Beaulieu, J.-P., et al. 2012, *Natur*, **481**, 167
 Chatterjee, S., Ford, E. B., Matsumura, S., & Rasio, F. A. 2008, *ApJ*, **686**, 580
 Chauvin, G., Lagrange, A.-M., Dumas, C., et al. 2004, *A&A*, **425**, L29
 Chauvin, G., Vigan, A., Bonnefoy, M., et al. 2015, *A&A*, **573**, 127
 Choi, J., Dotter, A., Conroy, C., et al. 2016, *ApJ*, **823**, 102
 Chung, S. J., Zhu, W., Udalski, A., et al. 2017, *ApJ*, **838**, 154
 Clanton, C., & Gaudi, B. S. 2014a, *ApJ*, **791**, 90
 Clanton, C., & Gaudi, B. S. 2014b, *ApJ*, **791**, 91
 Clanton, C., & Gaudi, B. S. 2016, *ApJ*, **819**, 125
 Clarkson, W., Sahu, K., Anderson, J., et al. 2008, *ApJ*, **684**, 1110

- Committee for a Decadal Survey of Astronomy & Astrophysics 2010, New Worlds, New Horizons in Astronomy and Astrophysics (Washington, DC: The National Academies Press)
- Cumming, A., Butler, R. P., Marcy, G. W., et al. 2008, *PASP*, **120**, 531
- Czekaj, M. A., Robin, A. C., Figueras, F., Luri, X., & Haywood, M. 2014, *A&A*, **564**, 102
- Dalcanton, J. J., Williams, B. F., Lang, D., et al. 2012, *ApJS*, **200**, 18
- Davenport, J. R. A., Becker, A. C., Kowalski, A. F., et al. 2012, *ApJ*, **748**, 58
- Dominik, M. 1998, *A&A*, **333**, L79
- Dong, S., Gould, A., Udalski, A., et al. 2009, *ApJ*, **695**, 970
- Doyle, L. R., Carter, J. A., Fabrycky, D. C., et al. 2011, *Sci*, **333**, 1602
- Dressler, A., Spergel, D., Mountain, M., et al. 2012, arXiv:1210.7809
- Dwek, E., Arendt, R. G., Hauser, M. G., et al. 1995, *ApJ*, **445**, 716
- Einasto, J. 1979, in IAU Symp. 84, The Large-scale Characteristics of the Galaxy, ed. W. B. Burton (Dordrecht: Reidel), 451
- Fressin, F., Torres, G., Charbonneau, D., et al. 2013, *ApJ*, **766**, 81
- Fulton, B. J., Petigura, E. A., Howard, A. W., et al. 2017, *AJ*, **154**, 109
- Gaudi, B. S. 1998, *ApJ*, **506**, 533
- Gaudi, B. S. 2012, *ARA&A*, **50**, 411
- Gaudi, B. S., Bennett, D. P., Udalski, A., et al. 2008, *Sci*, **319**, 927
- Gaudi, B. S., Naber, R. M., & Sackett, P. D. 1998, *ApJL*, **502**, L33
- Gaudi, B. S., Stassun, K. G., Collins, K. A., et al. 2017, *Natur*, **546**, 514
- Gehrels, N. 2010, arXiv:1008.4936
- Gilliland, R. L., Chaplin, W. J., Dunham, E. W., et al. 2015, *AJ*, **150**, 133
- Gilliland, R. L., Chaplin, W. J., Jenkins, J. M., Ramsey, L. W., & Smith, J. C. 2011, *ApJS*, **197**, 6
- Gillon, M., Jehin, E., Lederer, S. M., et al. 2016, *Natur*, **533**, 221
- Girardi, L., Groenewegen, M. A. T., Hatziminaoglou, E., & da Costa, L. 2005, *A&A*, **436**, 895
- Goldreich, P., & Tremaine, S. 1980, *ApJ*, **241**, 425
- Goldreich, P., & Ward, W. R. 1973, *ApJ*, **183**, 1051
- Gonzalez, O. A., Rejkuba, M., Zoccali, M., et al. 2012, *A&A*, **543**, A13
- Gould, A. 2008, *ApJ*, **681**, 1593
- Gould, A. 2009, arXiv:0902.2211
- Gould, A., Dong, S., Gaudi, B. S., et al. 2010, *ApJ*, **720**, 1073
- Gould, A., & Gaucheron, C. 1997, *ApJ*, **477**, 580
- Gould, A., Huber, D., Penny, M., & Stello, D. 2014a, *JKAS*, **48**, 93
- Gould, A., & Loeb, A. 1992, *ApJ*, **396**, 104
- Gould, A., Udalski, A., Shin, I.-G., et al. 2014b, *Sci*, **345**, 46
- Gould, A., & Yee, J. C. 2014, *ApJ*, **784**, 64
- Goulding, N. T., Barnes, J. R., Pinfield, D. J., et al. 2012, *MNRAS*, **427**, 3358
- Green, J., Schechter, P., Baltay, C., et al. 2011, arXiv:1108.1374
- Green, J., Schechter, P., Baltay, C., et al. 2012, arXiv:1208.4012
- Han, C. 2006, *ApJ*, **638**, 1080
- Hauschildt, P. H., Allard, F., Ferguson, J., Baron, E., & Alexander, D. R. 1999, *ApJ*, **525**, 871
- Hayashi, C. 1981, *PTSPS*, **70**, 35
- Henderson, C. B. 2015, *ApJ*, **800**, 58
- Henderson, C. B., Gaudi, B. S., Han, C., et al. 2014a, *ApJ*, **794**, 52
- Henderson, C. B., Park, H., Sumi, T., et al. 2014b, *ApJ*, **794**, 71
- Henderson, C. B., Poleski, R., Penny, M., et al. 2016, *PASP*, **128**, 124401
- Henderson, C. B., & Shvartzvald, Y. 2016, *AJ*, **152**, 96
- Holtzman, J. A., Watson, A. M., Baum, W. A., et al. 1998, *AJ*, **115**, 1946
- Hosek, M. W., Lu, J. R., Anderson, J., et al. 2018, *ApJ*, **855**, 13
- Howard, A. W., Marcy, G. W., Johnson, J. A., et al. 2010, *Sci*, **330**, 653
- Hunter, J. D. 2007, *CSE*, **9**, 90
- Ida, S., & Lin, D. N. C. 2004, *ApJ*, **616**, 567
- Johnson, J. A., Aller, K. M., Howard, A. W., & Crepp, J. R. 2010, *PASP*, **122**, 905
- Jones, E., Oliphant, T. E., Peterson, P., et al. 2001, SciPy: Open Source Scientific Tools for Python, <http://www.scipy.org/>
- Jurić, M., & Tremaine, S. 2008, *ApJ*, **686**, 603
- Kasting, J. F., Whitmire, D. P., & Reynolds, R. T. 1993, *Icar*, **101**, 108
- Kayser, R., Refsdal, S., & Stabell, R. 1986, *A&A*, **166**, 36
- Kennedy, G. M., Kenyon, S. J., & Bromley, B. C. 2006, *ApJL*, **650**, L139
- Kerins, E., Robin, A. C., & Marshall, D. J. 2009, *MNRAS*, **396**, 1202
- Kim, S.-L., Lee, C.-U., Park, B.-G., et al. 2016, *JKAS*, **49**, 37
- Kokubo, E., & Ida, S. 2002, *ApJ*, **581**, 666
- Kopparapu, R. K., Ramirez, R., Kasting, J. F., et al. 2013, *ApJ*, **765**, 131
- Koshimoto, N., Udalski, A., Beaulieu, J. P., et al. 2017, *AJ*, **153**, 1
- Kroupa, P. 2001, *MNRAS*, **322**, 231
- Kuijken, K., & Rich, R. M. 2002, *AJ*, **124**, 2054
- Lagarde, N., Robin, A. C., Reylé, C., & Nasello, G. 2017, *A&A*, **601**, 27
- Lambrechts, M., & Johansen, A. 2012, *A&A*, **544**, A32
- Latham, D. W., Mazeh, T., Stefanik, R. P., Mayor, M., & Burki, G. 1989, *Natur*, **339**, 38
- Laughlin, G., Bodenheimer, P., & Adams, F. C. 2004, *ApJL*, **612**, 73
- Leinert, C., Bowyer, S., Haikala, L. K., et al. 1998, *A&AS*, **127**, 1
- Liebig, C., & Wambsgans, J. 2010, *A&A*, **520**, A68
- Lin, D. N. C., Bodenheimer, P., & Richardson, D. C. 1996, *Natur*, **380**, 606
- Lissauer, J. J., Ragozzine, D., Fabrycky, D. C., et al. 2011, *ApJS*, **197**, 8
- Luhn, J. K., Penny, M. T., & Gaudi, B. S. 2016, *ApJ*, **827**, 61
- Majewski, S. R., Zasowski, G., & Nidever, D. L. 2011, *ApJ*, **739**, 25
- Mao, S., & Paczynski, B. 1991, *ApJL*, **374**, L37
- Marois, C., Macintosh, B., Barman, T., et al. 2008, *Sci*, **322**, 1348
- Marois, C., Zuckerman, B., Konopacky, Q. M., Macintosh, B., & Barman, T. 2010, *Natur*, **468**, 1080
- Marshall, D. J., Robin, A. C., Reylé, C., Schultheis, M., & Picaud, S. 2006, *A&A*, **453**, 635
- Mayor, M., & Queloz, D. 1995, *Natur*, **378**, 355
- Mennesson, B., Gaudi, S., Seager, S., et al. 2016, *Proc. SPIE*, **9904**, 99040L
- Mizuno, H. 1980, *PTPh*, **64**, 544
- Montet, B. T., Crepp, J. R., Johnson, J. A., Howard, A. W., & Marcy, G. W. 2014, *ApJ*, **781**, 28
- Murta, A. 2015, GPC: General Polygon Clipper library, Astrophysics Source Code Library, ascl:1512.006
- Naef, D., Latham, D. W., Mayor, M., et al. 2001, *A&A*, **375**, L27
- Nataf, D. M., Gould, A., Fouqué, P., et al. 2013, *ApJ*, **769**, 88
- Navarro, M. G., Minniti, D., & Ramos, R. C. 2017, *ApJL*, **851**, 13
- Nemiroff, R. J., & Wickramasinghe, W. A. D. T. 1994, *ApJL*, **424**, L21
- Nielsen, E. L., & Close, L. M. 2010, *ApJ*, **717**, 878
- Ofek, E. O. 2014, MATLAB package for astronomy and astrophysics, Astrophysics Source Code Library, ascl:1407.005
- Oke, J. B., & Gunn, J. E. 1983, *ApJ*, **266**, 713
- Oliphant, T. E. 2006, A Guide to NumPy, <http://web.mit.edu/dvp/Public/numpybook.pdf>
- Owen, J. E., & Wu, Y. 2017, *ApJ*, **847**, 29
- Paczynski, B. 1986, *ApJ*, **304**, 1
- Papaloizou, J. C. B., & Terquem, C. 2001, *MNRAS*, **325**, 221
- Pasetto, S., Grebel, E. K., Chiosi, C., et al. 2018, *ApJ*, **860**, 120
- Peale, S. 1997, *Icar*, **127**, 269
- Pejcha, O., & Heyrovský, D. 2009, *ApJ*, **690**, 1772
- Penny, M. T. 2014, *ApJ*, **790**, 142
- Penny, M. T., Kerins, E., Rattenbury, N., et al. 2013, *MNRAS*, **434**, 2
- Penny, M. T., Mao, S., & Kerins, E. 2011, *MNRAS*, **412**, 607
- Penny, M. T., Rattenbury, N. J., Gaudi, B. S., & Kerins, E. 2017, *AJ*, **153**, 161
- Perrin, M. D., Soummer, R., Elliott, E. M., Lallo, M. D., & Sivaramakrishnan, A. 2012, *Proc. SPIE*, **8442**, 84423D
- Perryman, M., Hartman, J., Bakos, G. Á., & Lindegren, L. 2014, *ApJ*, **797**, 14
- Petigura, E. A., Marcy, G. W., Winn, J. N., et al. 2018, *AJ*, **155**, 89
- Poleski, R., Skowron, J., Udalski, A., et al. 2014, *ApJ*, **795**, 42
- Poleski, R., Udalski, A., Bond, I. A., et al. 2017, *A&A*, **604**, 103
- Poleski, R., Zhu, W., Christie, G. W., et al. 2016, *ApJ*, **823**, 63
- Pollack, J. B., Hubickyj, O., Bodenheimer, P., et al. 1996, *Icar*, **124**, 62
- Ramirez, R. M., & Kallenecker, L. 2017, *ApJL*, **837**, L4
- Rasio, F. A., & Ford, E. B. 1996, *Sci*, **274**, 954
- Rattenbury, N. J., Mao, S., Debattista, V. P., et al. 2007, *MNRAS*, **378**, 1165
- Rauscher, B. J. 2015, *PASP*, **127**, 1144
- Rauscher, B. J., Fox, O., Ferruit, P., et al. 2007, *PASP*, **119**, 768
- Rivera, E. J., Lissauer, J. J., Butler, R. P., et al. 2005, *ApJ*, **634**, 625
- Robin, A. C., Marshall, D. J., Schultheis, M., & Reylé, C. 2012, *A&A*, **538**, A106
- Robin, A. C., Reylé, C., Derrière, S., & Picaud, S. 2003, *A&A*, **409**, 523
- Safronov, V. S. 1972, in IAU Symp. 45, The Motion, Evolution of Orbits, and Origin of Comets, ed. G. A. Chebotarev, E. I. Kazimirschak-Polonskaia, & B. G. Marsden (Dordrecht: Reidel), 329
- Sahu, K. C., Casertano, S., Bond, H. E., et al. 2006, *Natur*, **443**, 534
- Saito, R. K., Hempel, M., Minniti, D., et al. 2012, *A&A*, **537**, A107
- Sako, T., Sekiguchi, T., Sasaki, M., et al. 2007, *ExA*, **22**, 51
- Schmidt, S. J., Hawley, S. L., West, A. A., et al. 2015, *AJ*, **149**, 158
- Schultheis, M., Zasowski, G., Allende Prieto, C., et al. 2014, *AJ*, **148**, 24
- Seager, S. 2013, *Sci*, **340**, 577
- Sharma, S., Bland-Hawthorn, J., Johnston, K. V., & Binney, J. 2011, *ApJ*, **730**, 3
- Shvartzvald, Y., Maoz, D., Udalski, A., et al. 2016, *MNRAS*, **457**, 4089
- Shvartzvald, Y., Novati, S. C., Gaudi, B. S., et al. 2018, *ApJL*, **857**, 8
- Shvartzvald, Y., Yee, J. C., Novati, S. C., et al. 2017, *ApJL*, **840**, 3
- Simion, I. T., Belokurov, V., Irwin, M., et al. 2017, *MNRAS*, **471**, 4323
- Spergel, D., Gehrels, N., Baltay, C., et al. 2015, arXiv:1503.03757
- Spergel, D., Gehrels, N., Breckinridge, J., et al. 2013, arXiv:1305.5422
- Stanek, K. Z., Mateo, M., Udalski, A., et al. 1994, *ApJL*, **429**, 73

- Stern, D., Bartlett, J. G., Brodwin, M., et al. 2010, arXiv:1008.3563
- Sumi, T., Bennett, D. P., Bond, I. A., et al. 2010, *ApJ*, **710**, 1641
- Sumi, T., Bennett, D. P., Bond, I. A., et al. 2013, *ApJ*, **778**, 150
- Sumi, T., Kamiya, K., Bennett, D. P., et al. 2011, *Natur*, **473**, 349
- Sumi, T., & Penny, M. T. 2016, *ApJ*, **827**, 139
- Suzuki, D., Bennett, D. P., Sumi, T., et al. 2016, *ApJ*, **833**, 145
- Szymański, M. K., Udalski, A., Soszyński, I., et al. 2011, *AcA*, **61**, 83
- The Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, 33
- Thompson, S. E., Coughlin, J. L., Hoffman, K., et al. 2018, *ApJS*, **235**, 38
- Tofflemire, B. M., Wisniewski, J. P., Kowalski, A. F., et al. 2012, *AJ*, **143**, 12
- Udalski, A., Ryu, Y.-H., Sajadian, S., et al. 2018, *AcA*, **68**, 1
- Udalski, A., Szymański, M. K., & Szymański, G. 2015a, *AcA*, **65**, 1
- Udalski, A., Yee, J. C., Gould, A., et al. 2015b, *ApJ*, **799**, 237
- Udry, S., & Santos, N. C. 2007, *ARA&A*, **45**, 397
- Vanhollebeke, E., Groenewegen, M. A. T., & Girardi, L. 2009, *A&A*, **498**, 95
- Wang, J., Mawet, D., Hu, R., et al. 2018, *JATIS*, **4**, 035001
- Wegg, C., & Gerhard, O. 2013, *MNRAS*, **435**, 1874
- Wegg, C., Gerhard, O., & Portail, M. 2017, *ApJL*, **843**, 5
- Winn, J. N., & Fabrycky, D. C. 2015, *ARA&A*, **53**, 409
- Witt, H. J., & Mao, S. 1994, *ApJ*, **430**, 505
- Wolszczan, A., & Frail, D. A. 1992, *Natur*, **355**, 145
- Wyrzykowski, Ł., Rynkiewicz, A. E., Skowron, J., et al. 2015, *ApJS*, **216**, 12
- Yee, J. C. 2015, *ApJL*, **814**, 11
- Yee, J. C., Albrow, M., Barry, R. K., et al. 2014, arXiv:1409.2759
- Yee, J. C., Hung, L.-W., Bond, I. A., et al. 2013, *ApJ*, **769**, 77
- Yee, J. C., Shvartzvald, Y., Gal-Yam, A., et al. 2012, *ApJ*, **755**, 102
- Yoo, J., DePoy, D. L., Gal-Yam, A., et al. 2004, *ApJ*, **603**, 139
- Zhu, W., Penny, M., Mao, S., Gould, A., & Gendron, R. 2014, *ApJ*, **788**, 73
- Zoccali, M., Cassisi, S., Frogel, J. A., et al. 1999, *ApJ*, **530**, 418