



Sensitivity Analyses of Exoplanet Occurrence Rates from Kepler and Gaia

Megan I. Shabram¹, Natalie Batalha², Susan E. Thompson³, Danley C. Hsu^{4,5,6,7}, Eric B. Ford^{4,5,6,7},
Jessie L. Christiansen⁸, Daniel Huber^{9,10,11,12}, Travis Berger⁹, Joseph Catanzarite¹, Benjamin E. Nelson¹³, Steve Bryson¹,

Ruslan Belikov¹, Chris Burke¹⁴, and Doug Caldwell¹

¹ NASA Ames Research Center, Moffett Field, CA 94035, USA; megan.i.shabram@nasa.gov

² Department of Astronomy & Astrophysics, University of California, Santa Cruz, 95064, USA; Natalie.Batalha@ucsc.edu

³ Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA

⁴ Department of Astronomy & Astrophysics, 525 Davey Laboratory, The Pennsylvania State University, University Park, PA 16802, USA

⁵ Center for Exoplanets and Habitable Worlds, 525 Davey Laboratory, The Pennsylvania State University, University Park, PA 16802, USA

⁶ Center for Astrostatistics, 525 Davey Laboratory, The Pennsylvania State University, University Park, PA 16802, USA

⁷ Institute for CyberScience, The Pennsylvania State University, USA

⁸ Caltech/IPAC-NASA Exoplanet Science Institute, 770 S. Wilson Avenue, Pasadena, CA 91106, USA

⁹ Institute for Astronomy, University of Hawaii, 2680 Woodlawn Drive, Honolulu, HI 96822, USA

¹⁰ Sydney Institute for Astronomy (SfA), School of Physics, University of Sydney, NSW 2006, Australia

¹¹ SETI Institute, 189 Bernardo Avenue, Mountain View, CA 94043, USA

¹² Stellar Astrophysics Centre, Department of Physics and Astronomy, Aarhus University, Ny Munkegade 120, DK-8000 Aarhus C, Denmark

¹³ Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA), Department of Physics and Astronomy, Northwestern University, 2145 Sheridan Road, Evanston, IL 60208, USA

¹⁴ MIT Kavli Institute for Astrophysics and Space Research, 77 Massachusetts Avenue, 37-241, Cambridge, MA 02139, USA

Received 2018 September 12; revised 2020 April 20; accepted 2020 May 5; published 2020 June 12

Abstract

We infer the number of planets per star as a function of orbital period and planet size using Kepler archival data products with updated stellar properties from the Gaia Data Release 2. Using hierarchical Bayesian modeling and Hamiltonian Monte Carlo, we incorporate planet radius uncertainties into an inhomogeneous Poisson point process model. We demonstrate that this model captures the general features of the outcome of the planet formation and evolution around GK stars and provides an infrastructure to use the Kepler results to constrain analytic planet distribution models. We report an increased mean and variance in the marginal posterior distributions for the number of planets per GK star when including planet radius measurement uncertainties. We estimate the number of planets per GK star between 0.75 and 2.5 R_{\oplus} and with orbital periods of 50–300 days to have a 68% credible interval of 0.49–0.77 and a posterior mean of 0.63. This posterior has a smaller mean and a larger variance than the occurrence rate calculated in this work and in Burke et al. for the same parameter space using the $Q1$ – $Q16$ (previous Kepler planet candidate and stellar catalog). We attribute the smaller mean to many of the instrumental false positives at longer orbital periods being removed from the DR25 catalog. We find that the accuracy and precision of our hierarchical Bayesian model posterior distributions are less sensitive to the total number of planets in the sample, and more so for the characteristics of the catalog completeness and reliability and the span of the planet parameter space.

Key words: Exoplanets (498); Exoplanet catalogs (488); Transit photometry (1709); Bayesian statistics (1900); Astrostatistics (1882)

1. Introduction

NASA’s Kepler mission was designed to yield an ensemble of planetary systems amenable to statistical analysis (Borucki et al. 2010; Jenkins et al. 2010; Koch et al. 2010). During its primary phase, Kepler stared nearly continuously at a single field for 4 yr, monitoring approximately 190,000 stars that are mostly on the main sequence (Batalha et al. 2010; Brown et al. 2011). Kepler’s goal was to look for signs of transiting exoplanets and ultimately determine the frequency of temperate, Earth-size planets around Sun-like stars. This process led to a survey catalog of planet candidates with well-characterized completeness and reliability (Bryson & Morton 2017; Burke & Catanzarite 2017; Christiansen 2017; Coughlin 2017; Mullally 2017). Furthermore, Burke et al. (2015) investigated systematics in the derived occurrence rates caused by assumptions about the pipeline sensitivity, characterized by

Christiansen et al. (2015). The characterization of the Kepler pipeline sensitivity is critical to robust occurrence rate studies and future work that utilizes the results from Kepler.

With approximately 2327 confirmed planets and 2244 planet candidates from the Kepler mission (Borucki et al. 2011a, 2011b; Batalha et al. 2013; Batalha 2014; Burke et al. 2014; Mullally et al. 2015; Rowe et al. 2015; Borucki 2016), scientists are working to incorporate planet formation and subsequent evolution theories that can explain both the configuration of our solar system and planetary systems that can be very different from our own, for example, systems with dwarf stars and bright giants (Dressing & Charbonneau 2015; Silva Aguirre et al. 2017), single and binary host stars (Doyle et al. 2011; Orosz et al. 2012a, 2012b; Welsh et al. 2012, 2015), the number of planets in a system (Fabrycky et al. 2014; Lissauer et al. 2014), planet mass and size (Weiss & Marcy 2014; Rogers 2015; Wolfgang et al. 2016; Carrera et al. 2018), and orbital characteristics (Van Eylen & Albrecht 2015; Shabram et al. 2016). However, large uncertainties in stellar properties translate into large uncertainties in individual planet properties (Huber et al. 2014; Berger et al. 2018; Fulton & Petigura 2018)



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

and can limit studies attempting to characterize the exoplanet population. Despite the large uncertainties, we are able to develop generative models (i.e., the statistical process that describes how the data are generated) that handle large measurement uncertainty and highly correlated uncertainty of some planet candidate parameters. Additionally, sources of bias can be naturally incorporated into statistically robust occurrence rate analyses (Youdin 2011; Foreman-Mackey et al. 2014; Burke et al. 2015; Hsu et al. 2018, 2019). These population analyses are becoming more tractable, enabling a better understanding of the physical and orbital properties of exoplanet systems on a broad scale.

Standard occurrence rate studies have largely ignored the radius uncertainty contribution from the planet’s host star (Catanzarite & Shao 2011; Howard et al. 2012; Dong & Zhu 2013; Dressing & Charbonneau 2013, 2015; Petigura et al. 2013a, 2013b; Farr et al. 2014, 2015; Mulders et al. 2015, 2018, 2019; Silburt et al. 2015; Fulton et al. 2017; Van Eylen et al. 2018). Burke et al. (2015) characterize terrestrial planet occurrence rates for the Kepler GK dwarf sample, also without the inclusion of planet radius measurement uncertainties. Mulders et al. (2018, 2019) use a forward model with the latest Kepler data products to characterize planetary systems around stars (in addition to the number of planets per stellar type). Fulton & Petigura (2018) have investigated the stellar mass dependence of the planet radius gap using Gaia-updated stellar mass, stellar radius, and planet sizes for the Kepler sample. However, Fulton & Petigura (2018) do not include the impact of planet radius uncertainties, accounting for survey completeness in an inverse detection efficiency method, a method shown to bias occurrence rates toward smaller values in Foreman-Mackey et al. (2014) and Hsu et al. (2018).

The Gaia Data Release 2 has now provided more precise stellar measurement uncertainties (Berger et al. 2018; Gaia Collaboration et al. 2018). Updates to the stellar properties in the Kepler sample now enable more robust hierarchical Bayesian occurrence rate posterior distributions. The contribution to occurrence rate estimates from uncertainty in planet radius can be included in occurrence rate estimates by using the uncertainty in the measured planet-to-star radius ratio from transit light-curve modeling. To get the planet radius, the planet-to-star radius ratio is simply multiplied by the assumed host star radius point estimate. This has been done in Hsu et al. (2018), an approximate Bayesian computation occurrence rate analysis for GK stars. Foreman-Mackey et al. (2014) consider the contribution to the planet radius uncertainties from the measured planet-to-star radius ratio and stellar radius uncertainties in their occurrence rate analysis for GK stars. However, they use a nonparametric Bayesian method that makes it difficult to interpret some population-level parameters for planet formation and subsequent evolution theories. Hsu et al. (2019) use approximate Bayesian computation to include the host star radius uncertainties and planet-to-star radius ratio uncertainties by incorporating additional Kepler data products to accurately characterize the efficiency of planets being recognized as “threshold crossing events” (TCEs). In addition, Hsu et al. (2019) include catalog reliability in the occurrence rate calculations, which is also another important parameter that can affect the results.

Including measurement uncertainties in the occurrence rate calculations is impactful for many reasons. When using the

Kepler catalog of planet candidates to constrain hierarchical Bayesian models, we are able to marginalize over noise when reporting posteriors of the number of planets per star. Including the measurement uncertainty is necessary to avoid a bias due to only using a histogram of mean values to infer population distributions. Furthermore, the inclusion of measurement uncertainties can allow better exploration of population-level parameters that describe planet formation and subsequent evolution relations.

In this work, we use Hamiltonian Monte Carlo (HMC; Neal 2012; Carpenter et al. 2017) to perform hierarchical Bayesian model calculations. The HMC method is the state of the art for sampling hierarchical Bayesian models. HMC uses a kinetic energy term, taking advantage of the gradient of the target density to efficiently sample from high-dimensional posteriors. For example, HMC can handle the inclusion of measurement uncertainties and many population-level parameters, for likelihood-based continuous distribution models. Furthermore, HMC provides advanced diagnostics to look for sources of numerical bias and other model pathologies characteristic of using Markov Chain Monte Carlo (MCMC) methods to perform hierarchical Bayesian model calculations. Thus, HMC is a powerful sampling method and very applicable for this work.

Here we employ a hierarchical Bayesian model in conjunction with an HMC sampler to infer planet occurrence rates while including the contribution from the planet host star radius uncertainty in the uncertainties in planet size. We demonstrate the use of standard and advanced diagnostics to assess the application of HMC for performing our hierarchical Bayesian model calculations. We use this statistical framework to demonstrate the impact of subtle differences in host star categorization and small differences in selected planet radii and orbital period across varied completeness and reliability parameter spaces.

In Section 2 we describe the observations and parameter space used in our investigations. In Section 3 we explain the statistical framework for this work. In Section 4 we explore the sensitivity of our occurrence rate methodology to small changes in the selected stars, reliability, and completeness, the number of planets, and uncertainties in planet size. In Section 5 we discuss our experimental design and future research. In Section 6 we summarize the conclusions of this work.

2. Observations

In Sections 2.1 through 2.3, we describe the various stellar cuts, planet parameter cuts, and the detection model used in this work. We use the cuts described below to explore the sensitivity of posterior estimates of occurrence rates from our statistical framework to subtle changes in the selected stars, selected planet parameters, the inclusion of radius measurement uncertainties, and updated stellar properties from Gaia.

2.1. Stars

We apply our model to three stellar catalogs with two sets of stellar cuts. A summary of the stellar cuts can be found in Table 1, and a summary of the catalogs used can be found in Table 2. The first set of stellar cuts (labeled “GK cuts ↑”) describes stellar cuts similar to those used in Burke et al. (2015) and Hsu et al. (2018) using the Q1–Q16 catalog release

Table 1
Summary of GK Star Classifications

<i>stars</i> (GK cuts) \uparrow	<i>stars</i> (GK cuts) \downarrow
$T_{\text{eff}}: 4200\text{--}6100\text{ K}$	$T_{\text{eff}}: 3900\text{--}6000\text{ K}$
$R_* < 1.15$	$R_* < 1.35$
$\log g > 4.0$	$\log g > 3.8$

Note. “GK cuts \uparrow ” are similar to the stellar parameter cuts used in the occurrence rate studies for the $Q1 - Q16$ Kepler planet candidate catalog release (Mullally et al. 2016). “GK cuts \downarrow ” are similar to the stellar parameter selection used in the SAG 13 analysis to compare occurrence rates across different teams.

(Mullally et al. 2016). The upward-pointing arrow indicates that this selection of GK stars has more stars compared to our second definition of GK stars, which we label “GK cuts \downarrow .” This second case contains fewer stars and is similar to the cuts used in the NASA Exoplanet Program’s Study Analysis Group 13¹⁵ (see Table 2). We choose these two selections to investigate how sensitive our results are to relatively small differences in the definition of the stellar category of interest and to explore how much power the data have to explore trends in stellar properties while using the state-of-the-art Kepler planet and star catalogs.

Before selecting the GK stars to be analyzed with the updated Gaia stellar properties, we start with a selection of FGK stars from the Gaia Data Release 2 (Gaia Collaboration et al. 2018; van Leeuwen et al. 2018) cross-matched to the Kepler DR25 stellar catalogs (Mathur et al. 2017). All stars included in these cross-matched catalogs use the available Gaia stellar properties. Initially, the cross-match between the Kepler and Gaia catalogs is based on position alone. For some Kepler targets, there are multiple Gaia targets that match positionally. To uniquely identify the source, we computed a delta magnitude and looked at its distribution. We use various quality cuts that further reduce our cross-matched sample. The motivation for these cuts is to choose a sample of stars where we are reasonably confident that each is near the FGK main sequence and is less likely to be impacted by sources of dilution. Both a maximum parallax uncertainty (10%) and the GAIA data quality flags are chosen so as to provide a cleaner sample. For instance, binary stellar companions can contribute to excess scatter about the astrometric model. We note that unlike Berger et al. (2018), no extinction corrections were applied. This results in a set of 78,005 Kepler target stars.¹⁶ These selection criteria are applied in the following order:

1. First, we remove all duplicate Gaia source ID rows (these duplicates also share Kepler IDs).
2. We make a cut where the difference for all cross-matched targets between the Gaia G mean magnitude and the Kepler magnitude (with bandpasses that have similar overall shape, range, and median) is within 1.5σ of the median. We chose this threshold for (Gaia G)–(Kepler mag) that prevents matching more than one Gaia target to our Kepler targets, thus preventing us from using stellar

¹⁵ https://exoplanets.nasa.gov/system/internal_resources/details/original/680_SAG13_closeout_8.3.17.pdf

¹⁶ A table listing the 78,005 targets with their KIC and Gaia IDs, parameters, and parameter uncertainties can be found at github.com/mshabram/PyStan-Kepler-Exoplanet-Populations/blob/master/Sensitivity-Analyses-of-Exoplanet-Occurrence-Rates-from-Kepler-and-Gaia/Data/q1q17_dr25_gaia_fgk.csv.

Table 2
Summary of Selected Stars from Various Stellar Catalogs

Catalog	# <i>stars</i> (GK cuts) \uparrow	# <i>stars</i> (GK cuts) \downarrow
a. $Q1 - Q16$	91,446	N/A
b. DR25	88,807	81,882
c. DR25 + Gaia	N/A	44,597

Note. We can compare results across disparate stellar catalogs using hierarchical Bayesian analysis.

properties associated with a background or foreground star rather than the intended Kepler target. We address the slight difference between the Gaia G and Kepler by using the median of the differences.

3. Following Evans (2018), we select an astrometric goodness of fit in the along-scan direction (GOF_AL) of less than 20 and an astrometric excess noise of less than 5, to exclude potential poorly resolved binaries or other problematic targets.
4. We include parallax quality cuts using the processing flag outputs of the module that calculates astrophysical parameters for the Gaia target stars. We selected only targets for which the Priam processing flags (A and B) are zero. This selects strictly positive parallax values, colors close to the standard locus, and parallax error less than 0.05 mas (Andrae et al. 2018; Lindegren et al. 2018). We note that the sky position of the target stars does not change much over the full Kepler field. We assume that occurrence rates do not depend on a star’s position in the galaxy, so the dependence of parallax error on sky position does not introduce significant bias. This would become important for assessing occurrence rates between disk and halo stars.
5. Sources with Kepler magnitude less than 16 are removed, and we apply a magnitude cut of $0.5 < G_{bp} - G_{rp} < 1.7$ (Lindegren et al. 2018). This color cut is more precise than using the temperature from the Kepler Input Catalog and more uniform than using temperatures from the DR25 stellar catalog, for selecting FGK stars.
6. Furthermore, we use a six-iteration quadratic fit of the color–luminosity relation for the main sequence with $\log_{10}(1.75)$ width to select FGK targets.

We summarize the stellar catalog versions investigated in this work in Table 2, and we report the number of selected stars for each case. Here “ $Q1 - Q16$ ” refers to the version of the Kepler star and planet catalogs release that precedes the “DR25” catalog release. We evaluate occurrence rates for the DR25 and “DR25 + Gaia” (a version that uses Gaia-updated stellar properties) catalogs with the “GK cuts \downarrow ” selections that were designated during the NASA Exoplanet Exploration Program Analysis Group Study Analysis Group 13 (SAG 13) working group meeting. In this work, we analyze the $Q1 - Q16$ planet candidate catalogs to benchmark our methods and results against the previous work of Burke et al. (2015) and Hsu et al. (2018). Therefore, we only consider the “GK cuts \uparrow ” case ($T_{\text{eff}}: 4200\text{--}6100\text{ K}$, $R_* < 1.15$, and $\log g > 4.0$) with the $Q1 - Q16$ planet and star catalog. By comparing the $Q1 - Q16$ planet candidate catalog occurrence rates to occurrence rates using the DR25 planet candidate catalog, we can see the impact on occurrence rates when many of the instrumental false

Table 3
Summary of Planet Size and Orbital Period Ranges

	$R_{p_{\min}} - R_{p_{\max}} (R_{\oplus})$	$P_{\text{orb}_{\min}} - P_{\text{orb}_{\max}} (\text{days})$
<i>planets</i> ↓	1.00–2.00	50–200
<i>planets</i> ↑	0.75–2.50	50–300

Note. We select these fairly complete orbital period and planet size ranges to facilitate comparison between catalogs and previous work, and we assess the sensitivity of occurrence rate posteriors to the choice of planet parameters when using our parametric hierarchical Bayesian model.

positives at longer orbital periods have been removed from the DR25 catalog. The vetting process and reliability characterization can be found in Thompson et al. (2018).

In Section 4.6 we compare occurrence rate posteriors using the catalogs described here to the catalog provided in Berger et al. (2018). The Berger et al. (2018) catalog updates host star radius values using values that were spectroscopically derived in the California-Kepler Survey (CKS; Fulton et al. 2017; Johnson et al. 2017; Petigura et al. 2017), as well as stars cooler than 4000 K from Gaidos et al. (2016). However, the full population of stars searched by Kepler has not been updated with spectroscopic follow-up at this time.

2.2. Planets

We choose two different cuts in planet parameters. First, we consider planets with sizes that range from 1 to $2 R_{\oplus}$ and orbital periods that range from 50 to 200 days, referred to as the “*planets*↓” case in Table 3 and hereafter. These cuts span a parameter space for GK stars that has a slightly higher average detection completeness than the second case we investigate. The second case we refer to as “*planets*↑,” which includes planets with sizes between 0.75 and $2.5 R_{\oplus}$ and orbital periods between 50 and 300 days. This case now contains less reliable planet candidates and has a larger variance in completeness values across the planet parameter space. In this case, the top left corner of the completeness grid (Thompson et al. 2018) near $P_{\text{orb}} = 50$ days and $R_p = 2.5 R_{\oplus}$ has a higher reliability and completeness, while the opposite corner near $P_{\text{orb}} = 300$ days and at $R_p = 0.75 R_{\oplus}$ has a lower reliability and completeness. The “*planets*↓” case is contained within the “*planets*↑” case and has overall less variance than the “*planets*↑” case. The detection completeness model is discussed further in Section 2.3. These cuts were chosen to compare to previous work and to assess how subtle differences in the completeness and reliability and in the ranges in planet parameter space can influence occurrence rate posteriors.

2.3. Detection Model

We employ the analytic pipeline completeness model described in Section 2 of Burke et al. (2015) to compare our results against previous catalogs and for sensitivity analysis. We precompute the completeness over a 61×57 (planet radius \times orbital period) grid. We approximate the completeness as constant within each bin using the value calculated for each bin center after dividing the planet radius range by 61 and the orbital period range by 57. For the gamma cumulative distribution function (CDF) coefficients (shape a , $scale$, and $size$) that describe the average detection efficiency of selected GK stars for our DR25 and DR25 + Gaia catalog analysis, we use $a = 30.87$, $size = 0$, and $scale = 0.271$, with a plateau

factor of 0.94 (Christiansen 2017; Thompson et al. 2018). These coefficients are derived using a gamma CDF that is fit to a detection efficiency model that includes vetting completeness. For our $Q1$ – $Q16$ analysis, we use $a = 4.65$, $size = 0$, and $scale = 0.98$ (Burke et al. 2015). We calculate transit durations assuming a circular orbit and use the mean stellar radius estimates. Figure 2 of Burke et al. (2015) shows the absolute difference between the analytic model used in this study and the higher-fidelity completeness model available as part of the DR25 occurrence rate data product release. Since differences are largest (a relative fraction of approximately 0.06) toward longer orbital periods, we focus our analysis on the parameter space of $P_{\text{orb}} < 300$ days for GK stars. This allows us to investigate a region of parameter space with relatively high reliability and completeness. We have not included a model for reliability in our analyses; however, we have restricted our analyses to shorter orbital periods where reliability is higher based on estimates from Thompson et al. (2018). Furthermore, preliminary results show that occurrence rate posteriors are not significantly influenced by reliability when planet orbital periods are less than 300 days. A discussion of the impact of the latest DR25 pipeline completeness and reliability products will be available in C. J. Burke et al. (2019, in preparation). Future studies will explore more vigorous treatments of including vetting efficiency and numerical pipeline completeness models. We discuss this further in Section 5.2.

3. Statistical Framework

We calculate occurrence rates using the inhomogeneous Poisson point process method with a parametric rate intensity as implemented in Burke et al. (2015), Youdin (2011), and Gregory & Lored (1992), now using HMC (Neal 2012; Carpenter et al. 2017) and including planet radius measurement uncertainties.¹⁷

3.1. The Hierarchical Bayesian Model

For this study, we parameterize the rate intensity function of an inhomogeneous Poisson point process as a power-law scaling of the planet radius and the orbital period. The inhomogeneous Poisson point process is a natural choice of the likelihood function for the occurrence of exoplanets per star, where each planet occurrence that is counted is very nearly independent of each other planet occurrence that is counted (ignoring multiple-planet systems).

The likelihood for our model is adopted from Burke et al. (2015) and Youdin (2011), now with the addition of Gaussian noise in planet size:

$$\mathcal{L} = \left[\prod_{i=1}^{N_l} f_i \right] \exp(-N_{\text{exp}}), \quad (1)$$

where N_l is the number of selected planets after the cuts in stellar parameters, planet radius, and orbital period have been applied. N_{exp} is the number of expected detections in all bins,

¹⁷ Code can be found at: github.com/mshabram/PyStan_Kepler_Exoplanet_Populations.

defined as

$$N_{\text{exp}} = F_0 \int_{P_{\min}}^{P_{\max}} \int_{R_{\min}}^{R_{\max}} \left[\sum_{j=1}^{N_*} n_j(P_{\text{orb}}, R_p) \right] \times \left(\frac{P_{\text{orb}}}{P_0} \right)^\beta \left(\frac{R_p}{R_0} \right)^\alpha dP_{\text{orb}} dR_p, \quad (2)$$

where n_j is the survey completeness (see Section 2.3 and references therein), which is a function of the planet radius R_p , orbital period P_{orb} , and stellar properties. The survey completeness is precomputed outside of our hierarchical Bayesian model as was done in Burke et al. (2015), and it also depends on the stellar mass, stellar radius, and semimajor axis. The hyperparameters in this hierarchical Bayesian model are α (the power-law index for the planet radius distribution), β (the power-law index for the orbital period distribution), and F_0 (the integrated number of planets per star); f_l is the number density from the power-law scaling of the planet occurrence rate evaluated over the list of detected planets. As we numerically simulate this likelihood function using HMC (see Section 3.2), each planet can take on values normally distributed around the true planet size $R_p^{(i)}$ (with N_l latent variables corresponding to the number of planets in the sample) and reported standard deviation $\sigma_{R_p^{(i)}}$ of the observed planet radius $R_{\text{Obs}}^{(i)}$:

$$R_{\text{Obs}}^{(i)} \sim \text{Normal}(R_p^{(i)}, \sigma_{R_p^{(i)}}): T[R_{p_{\min}}, R_{p_{\max}}]. \quad (3)$$

The convolution of the true planet sizes with their measurement uncertainties is truncated, $T[R_{p_{\min}}, R_{p_{\max}}]$, so that draws that are outside the selected planet range (either *planets*↑ or *planets*↓) are not considered when numerically simulating the integral of the likelihood. The truncation allows the data to be described as resulting from a data-generating process that only produces values within an interval. In this case values that are drawn below and/or above the specified interval are treated as not observed. This allows us to investigate how the choices in cuts impact the resultant occurrence rate distributions. Creating a model that allows for planets with mean radius values outside the selected parameter space to enter into the calculation of the posterior distribution for the selected range is beyond the scope of this paper. In this work, by definition, if the planet’s true value exists inside the selected range, it does not exist outside the selected range.

This hierarchical Bayesian model also ignores the constant multiplicative factors resulting in the survey completeness only entering the equations in the number of expected detections for all bins (Youdin 2011). We note that the uncertainties in planet size mean that when using our hierarchical Bayesian analysis, there is a nonzero number of planets that have a nonzero chance of occurring outside of the selected range while sampling from our likelihood function. This effect will need to be explored in the future by allowing the number of planets in the sample N_l to have flexibility, and to exclude planets where draws do not land inside the given range. Currently, our method assumes that all the planets selected have true values within the planet radius ranges specified. We reason that this effect would be important when stitching together occurrence

rate analysis for different planet radius ranges with our current parametric method.

3.2. Hamiltonian Monte Carlo

We use the Stan Bayesian statistical modeling software (Carpenter et al. 2017) to perform numerical calculations. We utilize the extensive Stan diagnostics to assess the convergence of our HMC simulations. We use uniform priors ranging from -5 to 5 for our hyperparameters α , β , and $\ln F_0$. We advance four chains for 1500 warm-up iterations followed by 1500 sampling iterations.

The tree depth is a configuration parameter of the No-U-Turn-Sampler used by Stan that can impact efficiency.¹⁸ We set the maximum tree depth to 10. We increase the maximum to 11, which roughly doubles the compute time. Each chain has an energy Bayesian fraction of missing information (E-BFMI) of approximately 0.8. A low E-BFMI (< 0.02) for a given chain implies a problem with the adaptation phase, and those chains likely did not explore the posterior distribution efficiently (Betancourt 2016).

We obtain Gelman–Rubin statistics \hat{R} of 1.0 for all parameters, as well as zero divergent transitions. Gelman–Rubin statistics are used to evaluate the variance within and between Markov chains. Large Gelman–Rubin statistics indicate possible nonconvergence. A Gelman–Rubin value close to 1 indicates no sign of nonconvergence from this particular statistical test. Divergent transitions are an indication that your posterior estimates are biased from numerical error. We obtain effective sample sizes (ESSs) of ~ 4600 for α , 6000 for β , ~ 4300 for $\ln F_0$, and 6000 for all the latent variables $R_p^{(i)}$. The ESS is a measure of how many draws from the Markov chain are effectively independent after the burn-in phase.

4. Results

In order to investigate the sensitivity of occurrence rate posteriors from Kepler data to small changes in the selected stars, reliability and completeness, the number of planets, and uncertainties in planet size, we perform fits over both the R_p and P_{orb} ranges (“*planets*↑” and “*planets*↓,” described in Table 3). The posterior distribution for all model parameters, including the power-law parameter estimates that describe the general features of the outcome of planet formation and subsequent evolution, can be found on the github repo for this project.¹⁹ We assess the occurrence rate posteriors when making subtle changes to the definition of GK stars (described in Table 1). Figure 1 shows kernel density estimates of marginal posterior distributions for the occurrence rate (i.e., the number of planets per GK star, F_0). The key labels read from top to bottom corresponding to curves going from the left to right. The *stars*↑ (*stars*↓) label means more (fewer) stars, the *planets*↑ (*planets*↓) label means more (fewer) planets, and the σ ↑ (σ ↓) label means with (without) measurement uncertainty in planet size. The dashed lines help indicate the occurrence rates calculated using the slightly warmer set of stars, *GK cuts* ↑ (i.e., *stars*↑), described in Table 1. The thicker lines help

¹⁸ A brief guide to Stan’s warnings can be found at <http://mc-stan.org/misc/warnings.html>.

¹⁹ github.com/mshabram/PyStan_Kepler_Exoplanet_Populations/tree/master/Sensitivity-Analyses-of-Exoplanet-Occurrence-Rates-from-Kepler-and-Gaia/posterior-distributions

indicate the inclusion of planet radius measurement uncertainties (i.e., σ_r). We will refer to this figure in Sections 4.2 through 4.5. Summary statistics for the occurrence rate posterior distributions can be found in Table 4, and two sample Kolmogorov–Smirnov (K-S) statistics for pairs of these occurrence rate posterior distributions can be found in the Appendix.

4.1. Sensitivity to Selections in Planet Radius and Orbital Period

We investigate the sensitivity of occurrence rate posteriors to the range of planet radii and orbital periods by comparing across the two ranges in planet radius and orbital periods described in Table 3. Our “*planets*↓” case contains approximately half the number of selected planets as our “*planets*↑” case and lies in a slightly higher average completeness space for GK stars of interest. The top panel of Figure 1 shows the marginal posterior for the number of planets per GK star for which the selected number of planets follows from the “*planets*↓” case. The middle and bottom panels of Figure 1 show occurrence rate posteriors for the number of planets per GK star when using the “*planets*↑” case (which includes the planets from the “*planets*↓” case).

When comparing these two clusters of marginal posteriors, we see that the “*planets*↓” curves have less variance than the marginal posteriors for the cluster of the “*planets*↑” cases, even though the “*planets*↓” case contains approximately half the number of selected planets. This could be in part due to the completeness and reliability varying more across the “*planets*↑” case (the larger planet parameter space box). For example, when comparing the completeness between the “*planets*↑” case and the “*planets*↓” case, parts of the larger box (“*planets*↑”) are in a higher-completeness and higher-reliability space (i.e., at $P_{\text{orb}} = 50$ days and at $R_p = 2.5 R_{\oplus}$), while another section is in a lower-reliability and lower-completeness space (i.e., at $P_{\text{orb}} = 300$ days and at $R_p = 0.75 R_{\oplus}$). Therefore, we attribute the larger variance for occurrence rate posteriors for the “*planets*↑” cases in part to (a) the larger variance in the detection efficiency across this parameter space and (b) the larger span in parameter space covered by the power-law rate intensity parameterization. Furthermore, although we expect the “*planets*↑” cases to have larger occurrence rates than the *planets*↓ cases (because we are probing a larger domain), the “*planets*↑” occurrence rate posteriors could be overestimated owing to the low, unaccounted-for reliability in the corner near $0.75 R_{\oplus}$ and 300 days.

4.2. Sensitivity to Selected Stars

Subtle differences in stellar cuts can impact the number of planets selected, where more (fewer) stars results in a smaller (larger) planet occurrence rate posterior mean. Table 5 shows that the “*DR25 GK cuts (stars)*↓” case has approximately 8% fewer selected stars than the “*DR25 GK cuts (stars)*↑” case. For these two selections of GK star cuts used in this study, see Table 1. We compare occurrence rates across these subtle differences in selected stars to first assess how sensitive our occurrence rate posteriors are to the choice of target stars. The difference in occurrence rates across subtle changes in stellar parameter cuts can be assessed by comparing the dashed green/black and the dashed blue/purple curve pairs in the top panel of Figure 1 and the dashed orange/pink and dashed red/cyan

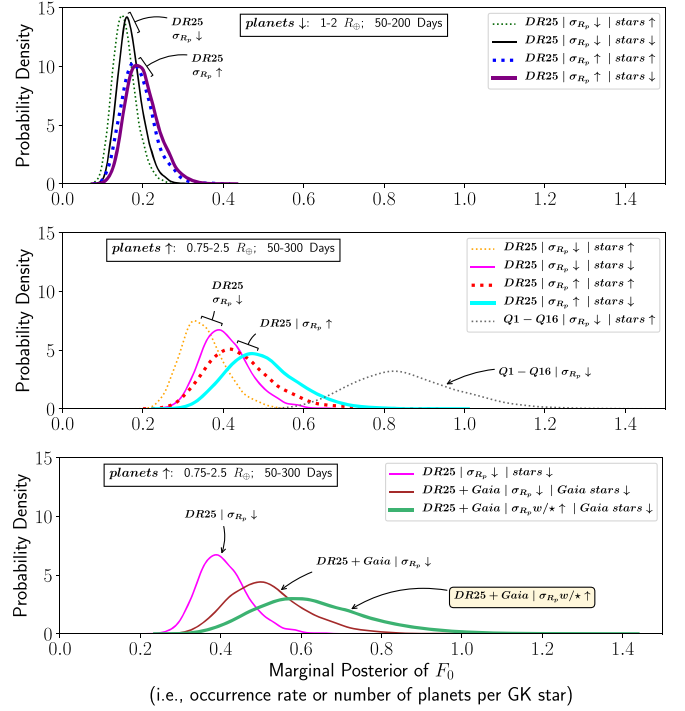


Figure 1. Kernel density estimates of marginal posterior distributions for the number of planets per GK star. The key labels read from top to bottom corresponding to curves going from the left to right. *stars*↑ (*stars*↓) means more (fewer) stars. *planets*↑ (*planets*↓) means more (fewer) planets. σ_r (σ_r ↓) means with (without) measurement uncertainty in planet size. The dashed lines help indicate the occurrence rates calculated using the slightly warmer set of stars, *GK cuts* ↑ (i.e., *stars*↑), described in Table 1. The thicker lines help indicate the inclusion of planet radius measurement uncertainties (i.e., σ_r). Excluding planet size (R_p) measurement uncertainty biases occurrence rates toward smaller values: compare dashed green/dashed blue and black/purple pairs in the top panel and dashed orange/dashed red and pink/cyan curve pairs in the middle panel. These correspond to fixed planet and star cuts with no measurement uncertainty/with measurement uncertainty (σ_{R_p} ↓/ σ_{R_p} ↑), respectively. A previous lower-reliability Kepler planet candidate catalog (Q1 – Q16 catalog) included more false positives, inflating the occurrence rate for this parameter space (dashed gray curve in the middle panel). Subtle differences in stellar cuts can impact the number of planets selected, where more stars result in fewer planets (compare dashed green/black and dashed blue/purple curves in the top panel and dashed orange/pink and dashed red/cyan curves in the middle panel). The occurrence rate variance is lower for planets in a slightly more complete part of parameter space (*planets*↓ in top panel) than in a slightly less complete part of parameter space (*planets*↑ in middle and bottom panels), even when there are fewer planets present in the *planets*↓ case. Although we expect the “*planets*↑” cases to have larger occurrence rates than the *planets*↓ cases (because we are probing a larger domain), the “*planets*↑” occurrence rate posteriors could be overestimated owing to the low, unaccounted-for reliability in the corner near $0.75 R_{\oplus}$ and 300 days. Comparing the pink and brown curves in the bottom panel shows the impact on occurrence rate posteriors in the *planets*↑ and *stars*↓ parameter space when using updated stellar radii from Gaia in the completeness model and updating planet sizes (and excluding measurement uncertainty). Propagating the stellar uncertainties from Gaia into the planet size (R_p) uncertainties while simultaneously updating stellar radii in the completeness model removes the bias toward smaller values and increases the variance of the occurrence rate (light-green curve in bottom panel).

cyan curve pairs in the middle panel of Figure 1. In these comparisons, the selected star parameters are varied while holding both the planet radius measurement uncertainty and the ranges in selected planet parameters fixed.

The “*planets*↓” case has a larger difference (8%) in the number of selected planets that make it through the two different GK stellar cut designations than the “*planets*↑,” yet this has a smaller influence on the difference in occurrence rate modes between these stellar cut designations. The “*planets*↑”

Table 4
Summary Statistics for Occurrence Rate Posterior Distributions

	Distribution	Mean	Var	Std	Mode	68% Credible Interval
<i>planets</i> ↓	<i>DR25</i> σ_{R_p} ↓ <i>stars</i> ↑	0.154	0.001	0.029	0.145	[0.126, 0.182]
	<i>DR25</i> σ_{R_p} ↓ <i>stars</i> ↓	0.168	0.001	0.03	0.157	[0.139, 0.198]
	<i>DR25</i> σ_{R_p} ↑ <i>stars</i> ↑	0.188	0.002	0.04	0.174	[0.149, 0.226]
	<i>DR25</i> σ_{R_p} ↑ <i>stars</i> ↓	0.199	0.002	0.04	0.179	[0.16, 0.239]
<i>planets</i> ↑	<i>DR25</i> σ_{R_p} ↓ <i>stars</i> ↑	0.35	0.003	0.055	0.322	[0.297, 0.405]
	<i>DR25</i> σ_{R_p} ↓ <i>stars</i> ↓	0.407	0.004	0.062	0.385	[0.348, 0.467]
	<i>DR25</i> σ_{R_p} ↑ <i>stars</i> ↑	0.442	0.007	0.085	0.4	[0.361, 0.523]
	<i>DR25</i> σ_{R_p} ↑ <i>stars</i> ↓	0.497	0.008	0.089	0.452	[0.411, 0.585]
	<i>Q1</i> – <i>Q16</i> σ_{R_p} ↓ <i>stars</i> ↑	0.854	0.017	0.131	0.812	[0.724, 0.985]
	<i>DR25</i> + <i>Gaia</i> σ_{R_p} ↓ <i>Gaia stars</i> ↓	0.519	0.01	0.098	0.483	[0.424, 0.613]
	<i>DR25</i> + <i>Gaia</i> σ_{R_p} w/★ ↑ <i>Gaia stars</i> ↓	0.63	0.02	0.141	0.606	[0.492, 0.767]

Note. We summarize the occurrence rate posterior distributions from Figures 1 and 2 via the mean, variance, standard deviation, mode, and 68.3% credible interval. The credible intervals are calculated such that the left- and right-hand regions of the posterior distribution outside the credible interval are equal in area.

case has a smaller difference (5%) in selected planets than the “*planets*↓” case and a larger difference in occurrence rate between occurrence rates calculated using these two stellar cut designations. The smaller difference in occurrence rate modes for the “*planets*↓” cases is likely in part due to the smaller area in parameter space, which must be described by the power-law rate intensity parameterization. Furthermore, differences in the occurrence rate posteriors between the two selections of GK stars may be from differences in the signal-to-noise regime (e.g., the *GK cuts* ↓ regime containing slightly larger maximum stellar radii and slightly cooler stars may let through more false positives in the “*planets*↑” case). It’s also possible that cooler GK stars host more planets, because we see the slight increase in occurrence rate posterior means in both selections of planet parameters. Mulders et al. (2015) find that the occurrence of Earth- to Neptune-sized planets is successively higher toward later spectral types at all orbital periods probed by Kepler.

4.3. Sensitivity to Planet Radius Measurement Uncertainties

Our analysis shows that when the planet-to-star radius ratio uncertainties are included, there is an upward shift in the occurrence rate posterior mean relative to when the planet-to-star radius ratio uncertainties are not included. In Figure 1, we can compare cases with fixed selected stars and planets for the *DR25* catalogs, including measurement uncertainties in planet size (indicated by “ σ_{R_p} ↑”) and not including them (indicated by “ σ_{R_p} ↓”). The “ σ_{R_p} ↓”/“ σ_{R_p} ↑” pairs of marginal posteriors for the number of planets per GK star are shown as dashed green/dashed blue and black/purple curve pairs in the top panel of Figure 1 and the dashed orange/dashed red and pink/cyan curve pairs in the middle panel of Figure 1, respectively.

The upward shift in the occurrence rate posterior mean can largely be attributed to (a) the wide range in uncertainty values across the planet radius sample. For planets that have well-constrained radius uncertainty, the location in completeness space stays relatively unchanged, whereas planets with large fractional radius uncertainties are more likely to have a large uncertainty in their detection completeness. (b) The detection probability is a sharp function of planet size near the detection threshold, with small planets more likely to be missed. For those small detected planets in the selected planet sample that

have larger planet radius measurement uncertainties, their observed radius will be more biased relative to their true radius. Both of those effects cause the model that ignores uncertainties to be biased toward a lower occurrence rate for more selected planets with radii near the threshold of detection.

4.4. Distribution Comparison to Burke+2015

We use a joint power-law rate intensity function in planet radius and orbital period for our inhomogeneous Poisson point process likelihood. This generative model is specified to capture broad features of the results of planet formation and subsequent evolution over small ranges in planet orbital period and radius. This likelihood and parameterization for Kepler exoplanet occurrence rates was put forth in Youdin (2011) and later applied by Burke et al. (2015), but neither of these studies included measurement uncertainties in a hierarchical Bayesian statistical framework. We re-create the conditions of Burke et al. (2015) to benchmark our methods and to evaluate how occurrence rates have changed when using the latest Kepler planet candidate catalog (the *DR25* planet candidate catalog). Our result for this occurrence rate is indicated as the dashed gray curve in the middle panel of Figure 1 and labeled “*Q1* – *Q16 GK cuts* ↑” for the “*planets*↑” case, without measurement uncertainties (“ σ_{R_p} ↓”), in the figure legend. In this case, we find an occurrence rate posterior mean of 0.85 with a 68% credible interval of 0.72–0.99 and an allowed range of 0.48–1.58. For the same set of stellar and planet parameter cuts, Burke et al. (2015) report an occurrence rate posterior mean of 0.77 with an allowed range of 0.3–1.9. We attribute the smaller posterior width and larger posterior mean calculated in this study to be from a combination of unaccounted-for differences in the custom catalog used in Burke et al. (2015) and the *Q1* – *Q16* catalog available at the NASA Exoplanet Archive, and potentially due to differences in the MCMC methods and diagnostics used.

4.5. Stars from Gaia

Using our statistical framework, we can compare disparate stellar catalogs. With the *Gaia*-updated stellar properties, the assumed stellar radii became larger on average (Berger et al. 2018). Additionally, the sample now has fewer evolved stars

for which Kepler has reduced planet detection efficiency due to their larger size. We first assess the impact of updated stellar radii from Gaia on the mean and variance of occurrence rate posteriors in this region of parameter space. Updating stellar radii with Gaia DR2 parameter estimates will change both the precomputed completeness functions and the resulting planet sizes. In this first case, we exclude the measurement uncertainties in planet size that come from both the uncertainty in R_p/R_* measurements and from stellar radius measurements. Thus, we simply multiply the planet-to-star radius ratios in the Kepler DR25 catalog by the new stellar radius estimates from Gaia DR2 and change the stellar radius estimates used in the completeness model. The resultant occurrence rate posterior is shown in the bottom panel of Figure 1 as the brown curve labeled “DR25 + Gaia (GK cuts ↓)” with “planets↑” cuts. Comparing this occurrence rate posterior to the pink curve in the middle and bottom panels of Figure 1 (“DR25 GK cuts ↓” with “planets↑” cuts) demonstrates the increase in the mean and variance of the occurrence rate posterior for the Gaia-updated planet radius point estimates and completeness inputs. The large difference in these two occurrence rates can be attributed to planets moving out of the planet radius range of interest, to changes in the precomputed completeness (due to shifting stellar radii values), and to planets and stars being removed from the sample when using more aggressive stellar cuts (described in Section 2.1).

Next, we propagate the stellar radius uncertainties from Gaia DR2 into the planet size (R_p) uncertainties (and are now included along with the contribution to the planet radius uncertainty from R_p/R_* measurements) while updating stellar radii in the precomputed completeness model. The resulting occurrence rate posterior is shown as the light-green curve in the bottom panel of Figure 1, exhibiting a much wider posterior (larger variance) than previous posteriors that did not include the contribution to the planet radius uncertainty due to the host star radius uncertainties. This marginal posterior has a 68% credible interval of 0.49–0.77 and a mean of 0.63. This occurrence rate posterior has a larger variance and a smaller posterior mean than the posterior for this parameter space using the Q1 – Q16 planet candidate catalog, which has a 68% credible interval of 0.72–0.99 and a posterior mean of 0.85, and larger than the occurrence rate posterior when using the DR25 planet candidate catalog alone, which has a 68% credible interval of 0.41–0.59 and a posterior mean of 0.50. This shows that previous studies have overestimated the occurrence rate in this region of parameter space, likely because previous lower-reliability Kepler planet candidate catalogs, such as the Q1 – Q16 catalog, likely included more false positives. However, selecting a cleaned stellar catalog partially compensates for this change.

4.6. Results Using the Berger+2018 Catalog

The Berger et al. (2018) study has provided a catalog of revised planet and star radius measurements using Kepler DR25 stars cross-matched with stars from Gaia DR2. Berger et al. (2018) use quality cuts similar to those described in Section 2.1 and also incorporate the stellar host star spectroscopic follow-up from the CKS (Fulton et al. 2017; Johnson et al. 2017; Petigura et al. 2017). Furthermore, the results from Berger et al. (2018) account for the impact of reddening. The orange curve in Figure 2 shows the marginal posteriors for the number of planets per GK star, over the *stars*↓ and *planets*↑




Figure 2. Kernel density estimates of marginal posterior distributions for the number of planets per GK star. The posteriors shown here are for planets with radii between 0.75 and 2.5 R_{\oplus} and orbital periods between 50 and 300 days. The brown curve corresponds to the occurrence rate posterior calculated using the DR25 Kepler star and planet catalogs with stellar radii updated by cross-matching with Gaia DR2 data and does not include measurement uncertainty in planet size. The green curve is this same case but now includes planet radius measurements uncertainties from the R_p/R_* measurements and from the uncertainties in stellar radius measurements when using Gaia data. The orange curve is the occurrence rate posterior when using the Berger et al. (2018) catalog that includes stellar parameters updated using spectroscopic follow-up for host stars only. The stellar sample in full is updated using Gaia DR2. The orange curve demonstrates that using heterogeneous stellar parameters introduces a large bias in occurrence rates.

parameter space, using the Berger et al. (2018) catalog. For this case, we find a 68% credible interval of 0.45–0.64 and a posterior mean of 0.55. This result is close to the result for the occurrence rate posterior distribution using the DR25 + Gaia cross-match (without updates using CKS) described in Section 2.1, when measurement uncertainties are not included (shown as the brown curve in Figures 1 and 2). The results when measurement uncertainties are included for the DR25 + Gaia catalogs (without spectroscopic host star follow-up) are shown as the green curve in Figures 1 and 2, for reference. Systematic differences in measurement uncertainties for stars with and without detected planets are not included in our statistical model. The orange curve shown in Figure 2 demonstrates that using the Berger et al. (2018) catalog that includes heterogeneous stellar parameters introduces a bias in the occurrence rate. We note that we have not included the impact of reddening in our catalog, which could impart differences in the base occurrence rate calculated before planet radius uncertainties are included. However, this would not account for the bias we see between the orange and green curves where planet radius uncertainties are included in the model. We note that this bias can be addressed with spectroscopic characterization of all stars in the catalog.

5. Discussion

The application of hierarchical Bayesian inference to infer planet occurrence rates handles a relatively small number of detected planets by pooling and mustering the strength of each constituent while learning about the population. By using HMC

Table 5
Summary of Selected Stars and Planets

Catalog	# planets↓/stars	# planets↑/stars
a. $Q1 - Q16(GK\ cuts)stars^\uparrow$	N/A	106/91,446(0.0012)
b. $DR25(GK\ cuts)stars^\uparrow$	54/88,807(0.0006)	118/88,807(0.0013)
b. $DR25(GK\ cuts)stars^\downarrow$	58/81,882(0.0007)	124/81,882(0.0015)
c. $DR25 + Gaia(GK\ cuts)stars^\downarrow$	N/A	85/44,597(0.0019)

Note. Subtle changes in stellar parameter selections can result in data sets with fewer stars having more planets. This effect is seen in occurrence rate posteriors, suggesting that our method may be sensitive to probing relations with stellar parameters, even when using simple planet formation and subsequent evolution distribution models. The differences in the number of selected stars have a negligible contribution to the expected number of planets per star due to the contribution from the completeness model used in the hierarchical Bayesian statistical framework. However, subtle differences in the number of selected planets could in part be due to unaccounted-for reliability.

to sample from our posterior, we can apply a high-dimensional hierarchical Bayesian model that has more parameters than measurements. There are more parameters than measurements because when we include planet radius measurement uncertainties (i.e., noise, error, and model assumptions), the planet radius for every planet in the model becomes a parameter in the model. By assessing how the occurrence rates behave in response to subtle difference in the inputs, we can see the positive impact of the Kepler science team’s efforts to provide high-quality occurrence rate data products, and we can evaluate the opportunities for advancing the depth of the science questions we are asking regarding exoplanetary systems.

Current analysis from Gaia data has provided stellar radii with average uncertainties of 8% (Berger et al. 2018). Our selected Gaia cross-matched stellar population has uncertainties of approximately 5% on average. This allows us to incorporate quality stellar data into the current occurrence rate framework we are using, parameterized by planet orbital period and planet radius.

Hierarchical parametric Bayesian exoplanet occurrence rate studies provide the foundation for constraining more complex exoplanet population distributions using Kepler data. As the data quality improves with complementary observations such as stellar follow-up, and with reprocessing of the current Kepler data using emerging statistical methods, scientists can begin to answer more in-depth questions in order to characterize planetary systems. In the following section we discuss occurrence rates from several angles: the population model, the data quality, the computational methods used to constrain hierarchical Bayesian models, and the science questions at hand.

5.1. Generative Model and Precomputing the Survey Completeness

The likelihood we use in this study assumes a rate intensity that is correlated between bins, similar to Burke et al. (2015) and Foreman-Mackey et al. (2014). This is important to consider when including planet radius measurement uncertainty in occurrence rate studies, since each planet’s size can now take on a variety of values. In this case, the data-generating process would be the outcome of planet formation and subsequent evolution, whereas a nonparametric Bayesian method such as a Gaussian Cox process would be agnostic to any planet formation and subsequent evolution relations.

In this initial study, we use a precomputed completeness grid over planet radius and orbital period described in Section 2.3. When assessing the impact on occurrence rates from planet radius measurement uncertainties, our precomputed completeness grid eases the computation. In order to include the contribution from

the host star radius into the planet radius uncertainty, we need to include the host star uncertainty into the calculation of the probability of detection, the geometric transit probability, and any functions in the completeness model that depend on stellar properties. In Section 4.5, we probe how occurrence rate posteriors change when using stellar properties from Gaia to update the stellar radius point estimates for each observed star, the means of the planet candidate radius measurement uncertainties, and the means of the host star radius measurement uncertainties. In this case we assess the impact of the contribution to the planet radius uncertainty from the host star radius uncertainty by approximating the completeness function as constant within each bin in planet radius and orbital period. We find that occurrence rate marginal posterior distributions are not changed when increasing the resolution of our completeness grid.

5.2. Future Work

Future studies to include the stellar radius uncertainties in the completeness model and therefore include the stellar radii as latent variables in our hierarchical Bayesian model may require the calculation of the completeness model in each iteration when sampling from the likelihood. This would replace the precomputed completeness we use in this study, which is used as input in our statistical framework. The analytic completeness models described in Burke et al. (2015) and used in this work take significantly less computational time than the numerical completeness functions available as part of the DR25 Kepler occurrence rate data products. Moving away from a precomputed completeness and using the latest numerical completeness models may require more advanced computing resources and techniques to constrain occurrence rate statistical frameworks that include stellar parameter measurement uncertainties.

Measurement uncertainties for orbital periods are negligible, but when reparameterizing in terms of insolation flux, uncertainties in stellar effective temperature, stellar multiplicity, stellar mass, and stellar radius could contribute significantly to the uncertainties in occurrence rates as a function of insolation flux. Updates to stellar effective temperatures from analysis of Gaia data will allow future studies to properly parameterize the occurrence rate in terms of insolation flux, as orbital distance is calculated from the stellar mass and orbital period, and the orbital distance estimate is used in the detection efficiency calculations. By including the completeness functions directly in the hierarchical Bayesian model’s data-generating process (instead of a precomputed completeness) in addition to a functional form for the planet formation and subsequent evolution model, it will be possible

to marginalize over uncertainties in stellar parameters. This will ultimately lead to constraining occurrence rates as a function of insolation flux (and other stellar parameters) in addition to planet parameters.

By using Gaia data to better constrain planet radius uncertainties and provide accurate fractional uncertainties for insolation flux, we can assess the impact of excluding measurement uncertainty in the occurrence rate parameterizations that go beyond the impact of the planet radius uncertainties investigated here. This will improve previous occurrence rates calculated in terms of insolation flux that are biased by the inverse detection efficiency method (e.g., Fulton & Petigura 2018). The large disparity in the number of selected stars for the different catalogs used to investigate changes in occurrence rate posteriors motivates including stellar parameter dependence directly in occurrence rate studies in the future.

When cross-matching the DR25 Kepler stellar catalog with the Gaia DR2 stellar parameters, we remove stars that have indications that they may be poorly resolved binaries. This provides results that are less contaminated with dilution from binarity than previous studies. Ciardi et al. (2015), Hirsch et al. (2017), Furlan et al. (2017, 2018), and Furlan & Howell (2017) have measured a nonnegligible planet radius correction factor to account for stellar multiplicity. Furthermore, Bouma et al. (2018) show that for terrestrial-sized planets, stellar multiplicity can contribute uncertainties in occurrence rates of approximately 50%. Stellar multiplicity is an important consideration for occurrence rates beyond the dilution of the planet radius by overestimating the size of its host star, as it can also impact the measured semimajor axis. In future studies, including a model of the impact of stellar binarity directly in the generative model used in this analysis will allow the impact of stellar binarity on occurrence rates to be measured.

Preliminary occurrence rate estimates of potentially habitable planets are lower with the new reliability estimates from the DR25 9.3 Kepler occurrence rate data products. This suggests that a vigorous treatment of the catalog reliability for occurrence rate studies will be necessary for learning about the population of potentially habitable planets.

By including planet radius measurement uncertainties in a parametric hierarchical Bayesian occurrence rate calculation, we have provided the foundation for researchers to use the Kepler data set to constrain parameters in analytic planet distribution models. This can be done by investigating these relations in place of the simplistic power-law intensity parameterization described in this work. Furthermore, Zink et al. (2019) show that the Kepler dichotomy can be filled in by accounting for the effects of multiplicity on the detection efficiency, and they provide improved estimates of the multiplicity distribution. Future studies can include this updated detection efficiency while also incorporating radius measurement uncertainties into the likelihood function.

6. Conclusion

When using our parametric hierarchical Bayesian model in conjunction with Gaia data to (i) remove stars that have indications they may be poorly resolved binaries, (ii) update the uncertainties in planet radii and in turn include the contribution of the host star radii in the uncertainty in planet radii, and (iii) update the stellar parameters in the completeness model,

- we estimate the GK star planet occurrence rate between 0.75 and $2.5 R_{\oplus}$ and 50–300 days to have a 68% credible interval of 0.49–0.77 and a mean of 0.63.

When using the Berger et al. (2018) catalog that includes spectroscopic follow-up of host stars only, Gaia-updated stellar radii, and reddening,

- we find that a bias is introduced into the occurrence rate posterior distributions when using heterogeneous stellar radius measurement uncertainties.

By performing a hierarchical Bayesian occurrence rate analysis in a particular part of planet parameter space with differences in reliability and completeness,

- we find an upward shift in the occurrence rate posterior mean and a larger posterior variance when including measurement uncertainty in planet radius.

When evaluating the sensitivity of planet occurrence rates to subtle changes in the selected stars,

- our results suggest that our hierarchical Bayesian models (Bayesian models that include measurement uncertainties) are less sensitive to subtle differences in stellar properties, and more so to the the selected ranges in planet parameters.

By evaluating a set of slightly cooler stars and a set of slightly warmer stars across two sets of selected planets with different completeness and reliability characteristics,

- we show that the choice of stellar cuts can influence the number of planet candidates selected over the planet radius and orbital period grid of interest.
- we find that the cooler star sample has a slightly higher occurrence rate posterior for both sets of selected planets.

This difference could in part be from (a) the slightly cooler selected stars letting through more false positives and (b) the fact that the slightly cooler set of GK stars could host more planets. Work by Dressing & Charbonneau (2013, 2015) found that cooler M dwarf stars have larger occurrence rates. Note that the Dressing & Charbonneau (2013, 2015) results are derived for much cooler stars (2600–4000 K) than those investigated here. This motivates the inclusion of a more vigorous treatment of the catalog reliability in future occurrence rate studies. Furthermore, the inclusion of stellar population-level parameters in hierarchical Bayesian occurrence rate studies will allow the characterization of the stellar dependence of exoplanet occurrence rates. It may be important to include the stellar dependence in statistically robust occurrence rate studies before we can select targets of opportunity for some exoplanet research.

We also evaluate the impact of selecting planets in a slightly higher average completeness space, compared to a part of parameter space with slightly less average and larger variance in completeness.

- We find that the selection of planets over the slightly more complete part of parameter space results in occurrence rate marginal posteriors with less variance than the space evaluated over a slightly less complete part of parameter space with more variance in completeness.

This is interesting because the “*planets*↓” case (slightly more complete space) contains approximately 50% fewer planets than the “*planets*↑” case.

- This suggests that the precision (variance) in the occurrence rate posteriors when using the statistical framework in this work is less sensitive to the number of planets that




Figure A1. Two-sample K-S statistics for occurrence rate posterior distribution comparisons.

make it through the planet cuts, and more so on (a) the span that the rate intensity parameterization is providing coverage over and (b) the effective number of stars searched. The effective number of stars searched (i.e., how efficient Kepler is at detecting planets in a given part of parameter space) depends on the characteristics of the completeness and reliability space and the signal-to-noise regime.

Shabram’s research is supported by an appointment to the NASA Postdoctoral Program at the NASA Ames Research Center, administered by Universities Space Research Association under contract with NASA. This manuscript benefited from discussions at the Stanford Stan User’s Group. This work benefited from Dan Foreman-Mackey’s blog post regarding an experiment in open science for Kepler exoplanet occurrence rates. This work makes use of the Stan open-source Bayesian modeling language, Jupyter, and Python software including the SciPy ecosystem and Seaborn. E.B.F. and D.C.H. acknowledge the Penn State Center for Astrostatistics and Center for Exoplanets and Habitable Worlds, which is supported by the Pennsylvania State University, the Eberly College of Science, and the Pennsylvania Space Grant Consortium. Computations for this research were

performed on the Pennsylvania State University’s Institute for CyberScience Advanced CyberInfrastructure (ICS-ACI). E.B.F. acknowledges support from NASA Exoplanets Research Program award No. NNX15AE21G. This work was partially supported by the NSF grant DMS 1127914 to the Statistical and Applied Mathematical Sciences Institute (SAMSI). The results reported herein benefitted from collaborations and/or information exchange within NASA’s Nexus for Exoplanet System Science (NExSS) research coordination network sponsored by NASA’s Science Mission Directorate. T.A.B. and D.H. acknowledge support by the National Science Foundation (AST-1717000) and the National Aeronautics and Space Administration under grants NNX14AB92G and NNX16AH45G.

Software: Stan (Carpenter et al. 2017), SciPy (Jones et al. 2001), Seaborn (Waskom et al. 2017), Jupyter (Kluyver et al. 2016).











Appendix

Two-sample Kolmogorov–Smirnov Statistics for Occurrence Rate Posterior Distribution Comparisons

In Figure A1, we show the “two-sample Kolmogorov–Smirnov (K-S) statistic” to assess the distance between pairs of

occurrence rate posterior distributions. A K-S statistic close to 0 means that the distributions are likely both drawn from the same underlying population, and a K-S statistic of 1 means that it is less likely that the distributions come from the same underlying distribution. The label colors correspond to the distribution plot color in Figures 1 and 2.

ORCID iDs

Megan I. Shabram  <https://orcid.org/0000-0003-1179-3125>
 Natalie Batalha  <https://orcid.org/0000-0002-7030-9519>
 Danley C. Hsu  <https://orcid.org/0000-0003-3447-1890>
 Eric B. Ford  <https://orcid.org/0000-0001-6545-639X>
 Jessie L. Christiansen  <https://orcid.org/0000-0002-8035-4778>
 Daniel Huber  <https://orcid.org/0000-0001-8832-4488>
 Travis Berger  <https://orcid.org/0000-0002-2580-3614>
 Steve Bryson  <https://orcid.org/0000-0003-0081-1797>
 Chris Burke  <https://orcid.org/0000-0002-7754-9486>
 Doug Caldwell  <https://orcid.org/0000-0003-1963-9616>

References

- Andrae, R., Fouesneau, M., Creevey, O., et al. 2018, *A&A*, **616**, A8
- Batalha, N. M. 2014, *PNAS*, **111**, 12647
- Batalha, N. M., Rowe, J. F., Bryson, S. T., et al. 2013, *ApJS*, **204**, 24
- Batalha, N. M., Rowe, J. F., Gilliland, R. L., et al. 2010, *ApJL*, **713**, L103
- Berger, T. A., Huber, D., Gaidos, E., & van Saders, J. L. 2018, *ApJ*, **866**, 99
- Betancourt, M. 2016, arXiv:1604.00695
- Borucki, W. J. 2016, *RPPH*, **79**, 036901
- Borucki, W. J., Koch, D., Basri, G., et al. 2010, *Sci*, **327**, 977
- Borucki, W. J., Koch, D. G., Basri, G., et al. 2011a, *ApJ*, **736**, 19
- Borucki, W. J., Koch, D. G., Batalha, N., et al. 2011b, AAS Meeting, **218**, 218.01
- Bouma, L. G., Masuda, K., & Winn, J. N. 2018, *AJ*, **155**, 244
- Brown, T. M., Latham, D. W., Everett, M. E., & Esquerdo, G. A. 2011, *AJ*, **142**, 112
- Bryson, S. T., & Morton, T. D. 2017, Planet Reliability Metrics: Astrophysical Positional Probabilities for Data Release 25, Kepler Science Document *KSCI-19108-001*
- Burke, C. J., Bryson, S. T., Mullally, F., et al. 2014, *ApJS*, **210**, 19
- Burke, C. J., & Catanzarite, J. 2017, Planet Detection Metrics: Per-Target Flux-Level Transit Injection Tests of TPS for Data Release 25, Kepler Science Document *KSCI-19109-002*
- Burke, C. J., Christiansen, J. L., Mullally, F., et al. 2015, *ApJ*, **809**, 8
- Carpenter, B., Gelman, A., Hoffman, M., et al. 2017, *Journal of Statistical Software*, **76**, 1
- Carrera, D., Ford, E. B., Izidoro, A., et al. 2018, *ApJ*, **886**, 104
- Catanzarite, J., & Shao, M. 2011, *ApJ*, **738**, 151
- Christiansen, J. L. 2017, Planet Detection Metrics: Pixel-Level Transit Injection Tests of Pipeline Detection Efficiency for Data Release 25, Kepler Science Document *KSCI-19110-001*
- Christiansen, J. L., Clarke, B. D., Burke, C. J., et al. 2015, *ApJ*, **810**, 95
- Ciardi, D. R., Beichman, C. A., Horch, E. P., & Howell, S. B. 2015, *ApJ*, **805**, 16
- Coughlin, J. L. 2017, Planet Detection Metrics: Robovetter Completeness and Effectiveness for Data Release 25, Kepler Science Document *KSCI-19114-001*
- Dong, S., & Zhu, Z. 2013, *ApJ*, **778**, 53
- Doyle, L. R., Carter, J. A., Fabrycky, D. C., et al. 2011, *Sci*, **333**, 1602
- Dressing, C. D., & Charbonneau, D. 2013, *ApJ*, **767**, 95
- Dressing, C. D., & Charbonneau, D. 2015, *ApJ*, **807**, 45
- Evans, D. F. 2018, *RNAAS*, **2**, 20
- Fabrycky, D. C., Lissauer, J. J., Ragozzine, D., et al. 2014, *ApJ*, **790**, 146
- Farr, W. M., Gair, J. R., Mandel, I., & Cutler, C. 2015, *PhRvD*, **91**, 023005
- Farr, W. M., Mandel, I., Aldridge, C., & Stroud, K. 2014, arXiv:1412.4849
- Foreman-Mackey, D., Hogg, D. W., & Morton, T. D. 2014, *ApJ*, **795**, 64
- Fulton, B. J., & Petigura, E. A. 2018, *AJ*, **156**, 264
- Fulton, B. J., Petigura, E. A., Howard, A. W., et al. 2017, *AJ*, **154**, 109
- Furlan, E., Ciardi, D. R., Cochran, W. D., et al. 2018, *ApJ*, **861**, 149
- Furlan, E., Ciardi, D. R., Everett, M. E., et al. 2017, *AJ*, **153**, 71
- Furlan, E., & Howell, S. B. 2017, *AJ*, **154**, 66
- Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2018, *A&A*, **616**, A1
- Gaidos, E., Mann, A. W., Kraus, A. L., & Ireland, M. 2016, *MNRAS*, **457**, 2877
- Gregory, P. C., & Lored, T. J. 1992, *ApJ*, **398**, 146
- Hirsch, L. A., Ciardi, D. R., Howard, A. W., et al. 2017, *AJ*, **153**, 117
- Howard, A. W., Marcy, G. W., Bryson, S. T., et al. 2012, *ApJS*, **201**, 15
- Hsu, D. C., Ford, E. B., Ragozzine, D., & Ashby, K. 2019, *AJ*, **158**, 109
- Hsu, D. C., Ford, E. B., Ragozzine, D., & Morehead, R. C. 2018, *AJ*, **155**, 205
- Huber, D., Silva Aguirre, V., Matthews, J. M., et al. 2014, *ApJS*, **211**, 2
- Jenkins, J. M., Caldwell, D. A., Chandrasekaran, H., et al. 2010, *ApJL*, **713**, L87
- Johnson, J. A., Petigura, E. A., Fulton, B. J., et al. 2017, *AJ*, **154**, 108
- Jones, E., Oliphant, T., Peterson, P., et al. 2001, SciPy: Open Source Scientific Tools for Python, <http://www.scipy.org/>
- Kluyver, T., Ragan-Kelley, B., Pérez, F., et al. 2016, in 20th International Conf. on Electronic Publishing, Positioning and Power in Academic Publishing: Players, Agents and Agendas, ed. F. Loizides & B. Schmidt (Amsterdam: IOS Press), 87
- Koch, D. G., Borucki, W. J., Basri, G., et al. 2010, *ApJL*, **713**, L79
- Lindgren, L., Hernandez, J., Bombrun, A., et al. 2018, *A&A*, **616**, A2
- Lissauer, J. J., Marcy, G. W., Bryson, S. T., et al. 2014, *ApJ*, **784**, 44
- Mathur, S., Huber, D., Batalha, N. M., et al. 2017, *ApJS*, **229**, 30
- Mulders, G. D., Mordasini, C., Pascucci, I., et al. 2019, *ApJ*, **887**, 157
- Mulders, G. D., Pascucci, I., & Apai, D. 2015, *ApJ*, **814**, 130
- Mulders, G. D., Pascucci, I., Apai, D., & Ciesla, F. J. 2018, *AJ*, **156**, 24
- Mullally, F. 2017, Planet Detection Metrics: Automatic Detection of Background Objects Using the Centroid Robovetter, Kepler Science Document *KSCI-19115-001*
- Mullally, F., Coughlin, J. L., Thompson, S. E., et al. 2015, *ApJS*, **217**, 31
- Mullally, F., Coughlin, J. L., Thompson, S. E., et al. 2016, *PASP*, **128**, 074502
- Neal, R. M. 2012, arXiv:1206.1901
- Orosz, J. A., Welsh, W. F., Carter, J. A., et al. 2012a, *ApJ*, **758**, 87
- Orosz, J. A., Welsh, W. F., Carter, J. A., et al. 2012b, *Sci*, **337**, 1511
- Petigura, E. A., Howard, A. W., & Marcy, G. W. 2013a, *PNAS*, **110**, 19273
- Petigura, E. A., Howard, A. W., Marcy, G. W., et al. 2017, *AJ*, **154**, 107
- Petigura, E. A., Marcy, G. W., & Howard, A. W. 2013b, *ApJ*, **770**, 69
- Rogers, L. A. 2015, *ApJ*, **801**, 41
- Rowe, J. F., Coughlin, J. L., Antoci, V., et al. 2015, *ApJS*, **217**, 16
- Shabram, M., Demory, B.-O., Cisewski, J., Ford, E. B., & Rogers, L. 2016, *ApJ*, **820**, 93
- Silburt, A., Gaidos, E., & Wu, Y. 2015, *ApJ*, **799**, 180
- Silva Aguirre, V., Lund, M. N., Antia, H. M., et al. 2017, *ApJ*, **835**, 173
- Thompson, S. E., Coughlin, J. L., Hoffman, K., et al. 2018, *ApJS*, **235**, 38
- Van Eylen, V., Agentoft, C., Lundkvist, M. S., et al. 2018, *MNRAS*, **479**, 4786
- Van Eylen, V., & Albrecht, S. 2015, *ApJ*, **808**, 126
- van Leeuwen, F., de Bruijne, J. H. J., Arenou, F., et al. 2018, Gaia DR2 Documentation, <https://gea.esac.esa.int/archive/documentation/GDR2/>
- Waskom, M., Botvinnik, O., O’Kane, D., et al. 2017, seaborn, v0.8.1, Zenodo, doi:10.5281/zenodo.883859
- Weiss, L. M., & Marcy, G. W. 2014, *ApJL*, **783**, L6
- Welsh, W. F., Orosz, J. A., Carter, J. A., et al. 2012, *Natur*, **481**, 475
- Welsh, W. F., Orosz, J. A., Short, D. R., et al. 2015, *ApJ*, **809**, 26
- Wolfgang, A., Rogers, L. A., & Ford, E. B. 2016, *ApJ*, **825**, 19
- Youdin, A. N. 2011, *ApJ*, **742**, 38
- Zink, J. K., Christiansen, J. L., & Hansen, B. M. S. 2019, *MNRAS*, **483**, 4479