

Scouting for dark showers at CMS and LHCb

Susan Born,¹ Rohith Karur^{2,3} , Simon Knapen,^{2,3} and Jessie Shelton^{1,4,5} 

¹*Department of Physics, University of Illinois Urbana-Champaign, Urbana, Illinois 61801, USA*

²*Theoretical Physics Group, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*

³*Berkeley Center for Theoretical Physics, Department of Physics, University of California, Berkeley, California 94720, USA*

⁴*Illinois Center for Advanced Studies of the Universe, University of Illinois Urbana-Champaign, Urbana, Illinois 61801, USA*

⁵*Center for Theoretical Physics, Department of Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*



(Received 3 April 2023; accepted 24 July 2023; published 23 August 2023)

We assess the capabilities of the CMS and LHCb searches for low- p_T displaced dimuon pairs to discover hidden valley models. To do so, we develop a new benchmark model featuring a light dark photon with dark flavor-violating couplings, which realizes a range of dimuon vertex topologies. We show that the data scouting techniques used in these searches provide unique sensitivity and we make some additional suggestions to further extend the scope of future experimental searches.

DOI: [10.1103/PhysRevD.108.035034](https://doi.org/10.1103/PhysRevD.108.035034)

I. INTRODUCTION

In its second decade of operation the Large Hadron Collider (LHC) is now transitioning to its high-intensity phase, which will deliver substantial new sensitivity for precision Standard Model (SM) measurements as well as new opportunities to search for very rare beyond-the-SM phenomena in the energy range from the GeV scale up to a few hundreds of GeV. In relation to the search for new physics, it has long been appreciated that trigger design can and should play a dynamical role in developing analyses, as opposed to a static obstruction when searching for soft signatures. All experiments therefore strive to perform as much event reconstruction and analysis as possible at the trigger level, as this allows for more informed decisions as to which events to commit to tape.

Because more and more of the event reconstruction can be performed in real time, ATLAS, CMS, and LHCb have also been able to record partial or reduced versions of the events reconstructed by their trigger-level algorithms. This allows for a much higher output rate, and thus greatly reduced trigger thresholds. This strategy is referred to as the turbo stream [1,2], data scouting [3], or trigger-level analysis [4] by LHCb, CMS, and ATLAS, respectively. At LHCb in particular, this online-analysis concept is a critical

component of their ongoing and future data taking strategies, with an $\mathcal{O}(1)$ fraction of all events already being recorded through the turbo stream [2]. This has enabled LHCb to measure (for example) certain exclusive charm decays and allowed them to perform searches for new, low-mass dimuon resonances [5–7]. Fully online strategies have also been deployed to great effect in searches at ATLAS and CMS, in particular for low-mass dijet [3,4] and (displaced) dimuon resonances [8–10]. A scouting search for low-mass diphoton resonances would furthermore be very well-motivated [11], though may be technically challenging. Broadly speaking, data scouting strategies will become even more powerful tools during the HL-LHC phase, provided that the existing capabilities can be maintained and, optimally, expanded, in the context of the higher instantaneous luminosity planned for future operation.

Data scouting is moreover of great interest from a theoretical point of view. The emerging appreciation that new physics may be relatively light and coupled only feebly to the SM (see e.g., the reviews [12,13]) places a renewed emphasis on searches for new phenomena in relatively low-energy final states. Data scouting techniques are very well-suited to making these challenging, high-rate kinematic regimes accessible. Moreover, searches relying on data scouting are inherently inclusive, insofar as e.g., a dimuon scouting trigger will record information about all dimuon pairs passing its selection criteria, independent of the properties of the rest of the event. This property makes data scouting searches particularly suitable to capture a wide range of unexpected phenomena. These two features are particularly important in the current and upcoming phases

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

of LHC operations, as theory priors are nowadays much weaker than those during the initial phase of the LHC.

Hidden valley (HV) models are a well-motivated and physically rich class of dark sector models [14], whose defining feature is the existence of interactions within their dark sector that generate a high-multiplicity final state. This is typically due to the presence of a confining gauge group in the dark sector that gives rise to dark showering and hadronization, in broad analogy to the QCD sector in the Standard Model (SM). Hidden valleys specifically have been developed as components of models seeking to address the origin of dark matter [15–19], the matter-antimatter asymmetry [20], and the stabilization of the electroweak scale [21–23]. More broadly, they represent a rather generic example of particle physics beyond the SM that could have flown under the radar of most existing searches at the LHC. Signatures of such hidden valleys often feature nonisolated objects, high-multiplicity and/or soft final states, and long-lived species, and thus frequently represent a challenging target for existing analysis strategies at the LHC. These “dark shower” signatures can give rise to spectacular and nonstandard final states, see e.g., [14,24–29], and are an emerging frontier for upcoming runs at the LHC [30].

A rigorous and comprehensive map of the theoretical possibilities within the dark shower framework is not yet possible. However, some useful mileage can be gained by systematically classifying and constraining the portals through which dark hadrons may decay back to the SM [31–34]. Notably, this strategy has demonstrated that models featuring (sub-)GeV-scale dark mesons can only give rise to a high multiplicity of visible final particles in a subset of possible portal scenarios. In particular, soft muonic final states are well-motivated, especially if the dark states are short- or moderately long-lived ($c\tau \lesssim 1$ cm). If the event is sufficiently energetic, these soft muonic final states can be searched for inside (semivisible) jets [35–39]. On the other hand, scouting techniques are the natural choice for scenarios where e.g., the H_T , dijet or missing transverse energy (MET) triggers are not very effective. The sensitivity of LHCb in particular to low-mass resonances produced in dark showers has been noted earlier [7,35,36,38], as well as potential additional sensitivity at Belle II [40]. In this work we take a close look at the complementarity of the recent CMS low-mass displaced dimuon search [9] with that at LHCb [7], and show that both are powerful probes of hidden valley models in which the dark particles have tracker-scale lifetimes. We also suggest a number of extensions of these analyses, which could extend their discovery potential to hidden valley models with different vertex topologies.

We carry out this analysis in the context of a benchmark model that has two flavors of light dark quarks and an elementary dark photon. This benchmark model has features that are reminiscent of both neutral naturalness models on one hand [21] and models of strongly interacting massive particle (SIMP) dark matter with a massive

dark photon on the other [41–44]. However, the primary aim of the model developed here is to serve as a useful reference for comparing and developing experimental sensitivities to dark shower events in the challenging low-mass regime. Our two-flavor dark photon model is designed to realize a variety of possible dimuon signatures with a relatively limited number of model parameters. It is thus proposed somewhat in the spirit of a simplified model approach to analysis design; however, we emphasize that it is a theoretically self-consistent, UV-complete model (as far as hadronization uncertainties allow).

The remainder of this paper is organized in two parts; the phenomenology and experimental sensitivity estimates are discussed in the bulk of the paper, while we reserve model-building aspects to the appendixes. In Sec. II we summarize the essential features of our benchmark model; theorists interested in the model itself may find it beneficial to read the appendixes here. We describe our procedure for recasting the existing CMS and LHCb searches in Sec. III and present results in Sec. IV. Our conclusions and recommendations are presented in Sec. V. Appendix A contains further theoretical details of the benchmark model and its PYTHIA implementation; some additional figures can be found in Appendix B.

II. BENCHMARK DARK SHOWER MODEL

Here we develop a relatively simple benchmark model that can realize a range of distinct dimuon signatures depending on the choices made for its mass spectrum. In this section we provide an overview of the model’s collider phenomenology, with further model details given in Appendix A.

A. Production

Throughout this study, we assume that the dark shower is initiated through an exotic decay of the SM Higgs boson. This is theoretically the simplest production mode, as it does not require the introduction of a new, heavy mediator to initiate the dark shower. The SM Higgs (h) is moreover very narrow, making it a sensitive probe of beyond the Standard Model physics in general [45,46] and hidden valley models in particular [47]. Studying dark shower production in exotic Higgs decays is especially well-motivated given the development of hidden valley models that aim to address the hierarchy problem, such as neutral naturalness models [21–23]. From an experimental point of view, Higgs production is also one of the more challenging possibilities owing to the relatively low- p_T final states that it generates. For our present purposes, this makes exotic Higgs decays an excellent benchmark scenario to illustrate the power of data scouting techniques, as alternative trigger paths for these events would need to rely on associated SM objects produced along with the Higgs in vector boson

fusion or associated VH production, with accordingly smaller production cross sections.

The Higgs boson initiates dark showers by decaying into pairs of dark quarks through an interaction of the form $h\bar{q}_D q_D$. This interaction can be readily constructed in the benchmark model described in detail in Appendix A 1; in this benchmark, moreover, the size of the dark quark couplings to the SM Higgs can be adjusted without significant impact on the dark hadron masses and lifetimes. For our purposes, the branching ratio $\text{Br}(h \rightarrow HV)$ into dark showers is the primary quantity of interest, and can be treated as a free parameter. We will therefore present our results in terms of this branching ratio.

B. Dark sector spectrum

Once the Higgs has decayed to a pair of dark quarks, those quarks will undergo showering and hadronization within the dark sector, which we model with the hidden valley module [48,49] in PYTHIA 8 [50]. The hidden valley module recently received a major update expanding its capabilities to describe dark meson multiplets with broken flavor symmetries [30], which was essential to implement the model we develop in this work.

We fix the number of colors and flavors to $N_c = 3$ and $N_f = 2$, respectively, and allow for the running of the dark gauge coupling. The two-flavor case is the most minimal setup that generates the range of displaced vertex topologies that we study in Sec. II C. We further add an additional, massive $U(1)$ interaction to facilitate the various decay chains described below. This $U(1)$ is spontaneously broken, and the resulting massive dark photon kinetically mixes with the SM photon to supply a muonphilic decay portal back to the SM. Figure 1 shows a schematic representation of the spectrum and the most important decay chains; the model is defined in full detail in Appendix A.

In a bit more detail, we assume that the confining dark gauge group dynamically breaks the hidden sector $SU(2) \times SU(2)$ flavor symmetry down to its diagonal subgroup, as is the case in the SM pion sector. This leads to three pseudo-Goldstone bosons, which we label π_1 , π_2 , and π_3 . There is also a heavier η meson, which corresponds to the would-be Goldstone boson of the anomalous axial $U(1)$. We further assume that the couplings of A' to dark sector quarks break parity, charge conjugation, and isospin. This isospin breaking will lead to small mass splittings between the various π_i , which are not relevant for the phenomenology we are interested in here. The charges of the A' are however chosen such that the following interaction is nonvanishing:

$$\mathcal{L} \supset g \sin \theta A'^\mu (\pi_2 \partial_\mu \eta - \eta \partial_\mu \pi_2), \quad (1)$$

where g is the $U(1)$ gauge coupling and $\sin \theta$ a mixing angle parametrizing the isospin violation. (See Appendix A 2




FIG. 1. Schematic representation of particle content and decay chains in the two-flavor hidden valley benchmark model developed in this work. The lowest-lying dark hadrons are an approximate dark-isospin triplet (π_1, π_2, π_3) and heavier pseudoscalar η . The \dots represent heavier mesons and baryons. The green arrow represents the dark parton shower and hadronization steps, which result in multiple η s and π_i s per event. Finally, the dark mesons couple to an elementary dark photon A' , which controls the (displaced) decays of the η and/or some of the π_i to final states with SM fermions, as detailed in Sec. II C.

for details). This interaction enables the decay $\eta \rightarrow A'\pi_2$. Both the π_3 and η can also decay to a pair of A' through the dark sector chiral anomaly, provided this channel is kinematically open. Meanwhile, the π_1 and π_2 are stable on collider timescales and contribute MET to the event.

Since the η is heavier than the π_i , one expects hidden sector hadronization to produce it less frequently. Hadronization, as always, is one of the major sources of theoretical uncertainty in hidden valley models. In particular, the relative probability of producing different particle species in dark hadronization cannot be calculated from first principles, and thus the choices made in the hadronization component of the PYTHIA HV module must also be considered as part of the definition of any benchmark model; for an up-to-date discussion of these choices see Ref. [30]. The PYTHIA cards for our model points are generated with our public PYTHON script [51] (more detail about the PYTHIA implementation can be found in Appendix A 4).

The available decay channels for the η are shown schematically in Fig. 2. If $m_\eta > 3m_\pi$, the prompt decay $\eta \rightarrow 3\pi$ will always dominate. This case leads to a marginally enhanced multiplicity of π 's, where the π_3 decays as $\pi_3 \rightarrow A'A' \rightarrow \text{SM}$. The $\eta \rightarrow A'A'$ decay mode is only open for $m_\eta > 2m_{A'}$, while $\eta \rightarrow A'\pi_2$ may occur if $m_\eta > m_{A'} + m_{\pi_2}$. If both channels are available (blue region in Fig. 2), the $\eta \rightarrow A'\pi_2$ channel dominates, especially if the η is long-lived. Finally, if none of these conditions are satisfied, the η can decay through an off shell A' into a SM final state plus π_2 . One expects this decay to occur with a proper decay length of $\mathcal{O}(\text{cm})$ or larger. We refer to Appendix A 3 for more details.




FIG. 2. Summary of the η phase space and available decay chains. The $\eta \rightarrow 3\pi$ decay (gray region) is always prompt, while the three-body decay $\eta \rightarrow \pi_2 \bar{f} f$ (red) tends to be displaced, with $f\bar{f}$ representing SM final states. The remaining modes can be either prompt or displaced, depending on the choice of model parameters. If the η is taken to be long-lived in the blue region, the $\eta \rightarrow A' \pi_2$ mode always dominates. The (A), (B₁), (B₂), and (C) labels refer to the decay topologies in Sec. II C.

C. Vertex topologies

Our benchmark model can generate three possible vertex topologies, as summarized in Fig. 3. Concretely, we consider the following scenarios:

- (i) *Scenario A (resonant, pointing)*: This scenario is realized when the $\eta \rightarrow 3\pi$ channel is open (gray region in Fig. 2) such that the π_3 meson is the only dark sector meson that decays directly to the SM. We assume that this occurs promptly, through the $\pi_3 \rightarrow A' A'$ channel, while the A' decays to the SM by kinetically mixing with the SM photon. If this mixing angle is small, A' can however be long-lived, and we treat its lifetime ($c\tau_{A'}$) as a free parameter. This scenario produces resonant dimuon pairs, for which the vector sum of their momenta points back to the beamline.




FIG. 3. Schematic representation of the decay topologies considered in this work, with the f (\bar{f}) SM (anti)fermions, e.g., μ^- (μ^+). The green lines indicate macroscopic decay lengths.

- (ii) *Scenario B (resonant, nonpointing)*: This case occurs when on shell dark photons are produced in the displaced decay of a long-lived parent particle. The vector sum of the momentum of the resulting resonant dimuon pairs does not point back to the beamline. We realize this in two different ways in our benchmark model:

Scenario B1 (higher visible multiplicity): In this case we keep the $\eta \rightarrow 3\pi$ and $\pi_3 \rightarrow A' A'$ decay channels open, but choose the dark gauge coupling small enough that the $\pi_3 \rightarrow A' A'$ decay becomes displaced (see Appendix A 2). For a given meson mass spectrum, this case has the same average multiplicity of displaced dimuon pairs as Scenario A. When the dark photon lifetime is taken to be prompt, the dimuon pairs are produced at the same vertex as the visible particles that come from the decay of the other dark photon.

Scenario B2 (lower visible multiplicity): In this scenario, we instead take the $\eta \rightarrow 3\pi$ and $\pi_3 \rightarrow A' A'$ channels to be kinematically closed, while $\eta \rightarrow \pi_2 A'$ is allowed (blue and green regions in Fig. 2). We further assume that the η is long-lived, but that the A' decays promptly to the SM, again by mixing with the SM photon. We treat the η lifetime ($c\tau_\eta$) as a free parameter. In this scenario, all three dark pions ($\pi_{1,2,3}$) now escape the detector, so the dimuon pairs are the only visible particles associated with the displaced vertex. In this case the average dimuon multiplicity depends on the average η multiplicity $\langle N_\eta \rangle$, and is substantially smaller than in Scenario A.

- (iii) *Scenario C (nonresonant)*: This scenario holds when the η can only decay through the three-body process involving an off-shell dark photon $\eta \rightarrow \pi_2 f \bar{f}$ (red region in Fig. 2). The η is typically long-lived in this case (see Appendix A 3). This scenario produces nonresonant dimuon pairs, with the same dependence on $\langle N_\eta \rangle$ as in Scenario B2.¹ In this scenario all dark pions ($\pi_{1,2,3}$) are again detector-stable.

As we will see, the existing CMS and LHCb searches already have excellent sensitivity for Scenario A, but the CMS search is not currently targeting Scenario B. Neither search currently targets Scenario C.

D. Parameter choices and hadron multiplicities

Our $N_f = 2$ confining dark sector has a substantial number of parameters, some of which are related to one another by chiral perturbation theory, as worked out in

¹Displaced vertices with nonresonant muon pairs can also be realized for $m_\eta > 3m_\pi$ and $2m_{A'} > m_{\pi_3}$, in which case π_3 has three-body decays $\pi_3 \rightarrow A' f \bar{f}$; however, in this case the π_3 proper lifetime tends to be prohibitively long in the kinematic regime relevant for scouting searches.

Appendixes A 1 and A 2. We choose the following parameters as our maximal independent set; the number of dark sector colors (N_c), the mass of π_2 (m_{π_2}), the mass of the η (m_η), the dark sector confinement scale (Λ), the scale of chiral symmetry breaking (f), the degree of isospin breaking ($\sin\theta$), the dark photon mass ($m_{A'}$), the dark $U(1)$ gauge coupling (g) and the mixing parameter of the A' with the SM photon (ϵ). The mass spectrum and decay tables of the dark sector π_i , η and A' are fully specified in terms of these parameters, as provided in the appendixes and by our PYTHON tool [51].

We follow PYTHIA's definition of the confinement scale Λ , where it is given by the dimensional transmutation parameter of the one-loop β -function of the dark sector $SU(N_c)$ coupling.² One may expect this scale to roughly correspond to the scale at which chiral perturbation theory ceases to be reliable, in other words

$$4\pi f \approx \Lambda. \quad (2)$$

For definitiveness, we assume this relation to hold exactly; changing this assumption would lead to $\mathcal{O}(1)$ variations in the decay widths of some of the dark mesons. In addition, our experience with the SM suggests that the meson associated with the anomalous axial $U(1)$ flavor symmetry has a mass comparable to the confinement scale. In our model, this role is played by the η -meson and as such we expect

$$m_\eta \approx \Lambda. \quad (3)$$

This expectation famously does not hold for theories in the large- N_c limit [52]. We will take Eq. (3) to hold as an equality as well and choose $N_c = 3$. We further take $\sin\theta = 0.1$, such that isospin breaking is relatively small and

$$m_{\pi_1} \approx m_{\pi_2} \approx m_{\pi_3} \quad (4)$$

for all our benchmark points. Choosing a small value of the isospin-breaking parameter $\sin\theta$ is also important for ensuring that the PYTHIA hidden valley module provides a reasonable approximation to the hadronization process, as we detail in Appendix A 4.

The parameters g and ϵ on the other hand are responsible for setting the lifetimes of the A' , π_3 , and η . As such, all our benchmark points in this paper only differ by different choices for the masses m_{π_3} , m_η , and $m_{A'}$, and the choices for the lifetimes of the corresponding particles ($c\tau_{\pi_3}$, $c\tau_\eta$, and $c\tau_{A'}$). Concretely, the individual scenarios are defined by choosing

²While this work was in the final stages of completion, PYTHIA version 8.309 was released, which now also includes the two- and three-loop $SU(N_c)$ β functions. All our results were obtained with PYTHIA version 8.308.

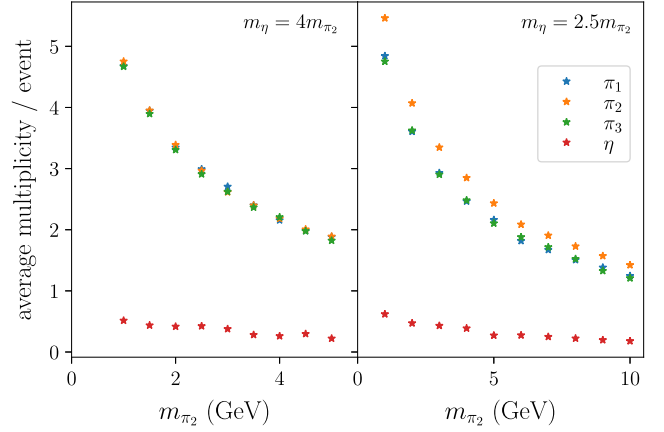


FIG. 4. Meson multiplicities for two representative choices of the m_η/m_{π_2} mass ratio. The left panel corresponds to Scenarios A and B1, while the right panel corresponds to Scenarios B2 and C.

- (i) *Scenario A*: $c\tau_{\pi_3} = c\tau_\eta = 0$, $m_\eta > 3m_{\pi_2}$
- (ii) *Scenario B1*: $c\tau_{A'} = c\tau_\eta = 0$, $m_\eta > 3m_{\pi_2}$,
- (iii) *Scenario B2*: $c\tau_{A'} = 0$, $m_\eta > m_{\pi_2} + m_{A'}$, π_3 detector-stable,
- (iv) *Scenario C*: $c\tau_{A'} = 0$, $m_\eta < m_{\pi_2} + m_{A'}$, π_3 detector-stable,

where “zero” $c\tau$ is meant as shorthand for “promptly decaying”, which is consistent with experimental constraints on visible dark photons for $m_{A'} \gtrsim 10$ MeV [53]. The remaining parameters are specified in the relevant figures.

With these assumptions, we can investigate average particle multiplicities, as generated by the PYTHIA 8 hidden valley module. We show the average meson multiplicities in Fig. 4 as a function of pion mass for two example choices of the mass ratio m_η/m_{π_2} . Naturally, the multiplicity of all meson species rises as the meson mass and thus the confinement scale are lowered. We furthermore see that the π_i multiplicity is a factor of several higher than the η multiplicity. This is in part due to the lower mass of the π_i , which means it is more likely to be produced in the hadronization model. Moreover, the dark-sector isospin-triplet vector mesons, the analogs of the ρ -mesons in the SM, are assumed to decay promptly to the π_i , further enhancing their multiplicity.

At the level of the meson mass spectrum, the difference between Scenarios (A, B1) vs Scenarios (B2, C) is whether the $\eta \rightarrow 3\pi$ channel is kinematically open. The former case comes with a slightly higher pion multiplicity. In Scenarios B2 and C, the $\eta \rightarrow A'\pi_2$ decay mode explicitly breaks isospin, and is the reason why the π_2 multiplicity is higher than the $\pi_{1,3}$ multiplicities in the right-hand panel of Fig. 4.

III. ANALYSIS

We consider two searches for low-mass displaced dimuon resonances from LHCb [7] and CMS [9]. Both

analyses are powerful probes of the dark shower signatures presented in the previous section. Here we briefly describe both searches, as well as our procedures for event generation and recasting.

A. CMS displaced dimuon search

The recent CMS search [9] records displaced low-mass dimuon pairs with p_T thresholds as low as ~ 3 GeV, at a rate of 3 kHz. This large rate is feasible because CMS only records a very small fraction of the total event information, which includes the muon four-momenta, number of hits per muon track, muon isolation, and track-quality information. The search used 101 fb^{-1} of data at 13 TeV.

The most important selection criteria in the CMS search are listed in Table I. These include cuts on the transverse momentum (p_T) and pseudorapidity (η) of the muons, as well as a cut on the distance between the dimuon vertex and the beamline (L_{xy}). In addition CMS also requires the vector sum of the muon momentum vectors to point back to the beamline in order to suppress backgrounds from fake vertices. Concretely, they impose a hard cut on the azimuthal angle between the vector sum of the muon momenta and the vector connecting the beamline to the displaced vertex. We will call this variable $\Delta\phi(\sum \vec{p}_T(\mu), \vec{x})$, with \vec{x} representing the location of the displaced vertex in the detector frame. This pointing requirement is automatically satisfied at truth level for Scenario A in Sec. II C, but not for Scenarios B and C. We will speculate on loosening this restriction in Sec. IV B. CMS also imposes further selection criteria, such as track and vertex quality requirements; we refer to [9] for more details. In addition to this set of baseline cuts, CMS defines a set of increasingly restrictive signal regions, the most relevant of which are listed in Table I. In particular, an event is put in the “isolated” category if the scalar sum of the p_T of all tracks within a cone of $\Delta R < 0.3$ around each muon does not exceed 20% of the p_T of the muon.

CMS provides trigger and offline efficiencies for each signal region, as well as the number of excluded events. We use this data in our reinterpretation, as described in Sec. III C. The collaboration evaluates signal efficiencies for two well-motivated benchmark topologies; an exotic B -meson and exotic Higgs decay, respectively $B \rightarrow X_s S$ and $h \rightarrow A'A'$, where the scalar (S) and vector (A') are assumed to decay to $\mu^+\mu^-$. We validated our reinterpretation procedure by simulating both signal models and verifying that for each model the exclusions we estimate by applying the provided efficiencies are in excellent agreement with the experimental results. Since we expect the muon p_T spectrum to be rather soft for most low-mass hidden valley models, we consider B -meson decay to be the more appropriate comparison model.

B. LHCb displaced dimuon search

The LHCb detector is specifically optimized to search for low-mass resonances and has consistently been on the forefront of the usage of online analysis strategies. This has resulted in excellent sensitivity to light beyond-the-SM particles, in particular if these particles have a substantial branching ratio to muons [6,7,54,55]. Here we focus on LHCb’s recent inclusive search for displaced dimuon resonances [7], which was carried out using 5.1 fb^{-1} of data at 13 TeV. The most important fiducial cuts are listed in Table I. The main differences with respect to the CMS selection are: (i) a substantially lower p_T cut on the muons; (ii) the more forward acceptance of the LHCb detector; (iii) the smaller L_{xy} range; and (iv) a loose cut on the opening angle between the two muons $\alpha(\mu^+, \mu^-)$.

C. Simulation framework

For our signal Monte Carlo samples, we generated Higgs bosons through the gluon fusion channel using PYTHIA 8.308 [50] and implemented the decay mode to dark quarks.

TABLE I. Most important fiducial cuts for the CMS [9] and LHCb [7] searches, as well as most relevant signal regions for the purpose of the study in this work. The CMS signal regions are nonexclusive, while the LHCb signal regions are exclusive.

	Fiducial cuts	Signal regions
CMS	$p_T(\mu) > 3 \text{ GeV}$	$0.0 \text{ cm} < L_{xy}$ (with and without isolation)
	$ \eta(\mu) < 2.4$	$0.2 \text{ cm} < L_{xy}$ (with and without isolation)
	$L_{xy} < 11 \text{ cm}$	$1.0 \text{ cm} < L_{xy}$ (with and without isolation)
	$\Delta\phi(\sum \vec{p}_T(\mu), \vec{x}) < 0.02$	$2.4 \text{ cm} < L_{xy}$ (with and without isolation)
		$3.0 \text{ cm} < L_{xy}$ (with and without isolation)
		$7.0 \text{ cm} < L_{xy}$ (with and without isolation)
LHCb	$p_T(\mu) > 0.5 \text{ GeV}$, $ \vec{p}(\mu) > 10 \text{ GeV}$	$2 \text{ GeV} < p_T(A') < 3 \text{ GeV}$ (with and without pointing)
	$2 < \eta(\mu) < 4.5$	$3 \text{ GeV} < p_T(A') < 5 \text{ GeV}$ (with and without pointing)
	$\sqrt{p_T(\mu^+)p_T(\mu^-)} > 1.5 \text{ GeV}$	$5 \text{ GeV} < p_T(A') < 10 \text{ GeV}$ (with and without pointing)
	$1.2 \text{ cm} < L_{xy} < 3 \text{ cm}$	
	$\alpha(\mu^+, \mu^-) > 3 \text{ mrad}$	

The dark quarks are subsequently showered and hadronized with the PYTHIA 8 hidden valley module [48,49]. Further detail on the PYTHIA implementation of our benchmark model is given in Appendix A 4. All signal samples were generated with promptly decaying particles and the lifetime dependence of the acceptance was accounted for by the reweighting procedure described in [56]. Concretely, for each event, we identify all vertices that pass the selection cuts in Table I, except for the L_{xy} requirements. For the CMS analysis, we moreover compute the isolation criteria on a vertex-by-vertex basis and classify events accordingly. We work with the truth-level four-momenta of the muons, effectively neglecting smearing due to the finite detector resolution. Reconstruction and trigger efficiencies were however incorporated to the extent possible, as described below.

To efficiently account for the cuts on L_{xy} , we assign an acceptance weight $w_{\text{acc}}(v)$ to each vertex v that passed these cuts,

$$w_{\text{acc}}(v) = e^{-\frac{L_{xy}^-}{\beta\gamma \cosh(\eta)c\tau}} - e^{-\frac{L_{xy}^+}{\beta\gamma \cosh(\eta)c\tau}}, \quad (5)$$

which corresponds to the probability that the long-lived particle corresponding to v decays within two coaxial cylinders with inner and outer radii L_{xy}^- and L_{xy}^+ . In other words, here the values $[L_{xy}^-, L_{xy}^+]$ represent the edges of the L_{xy} bins used in the analysis. Here $c\tau$, $\beta\gamma$, and η are respectively the proper lifetime, boost, and pseudorapidity of the long-lived particle. The full weight of the vertex is obtained after multiplying $w_{\text{acc}}(v)$ by the probabilities that the vertex would pass the trigger and offline selections, denoted by $\epsilon_{\text{trig}}(v)$ and $\epsilon_{\text{off}}(v)$, or

$$w(v) \equiv w_{\text{acc}}(v) \times \epsilon_{\text{trig}}(v) \times \epsilon_{\text{off}}(v). \quad (6)$$

For the CMS analysis, $\epsilon_{\text{trig}}(v)$ and $\epsilon_{\text{off}}(v)$ are provided in the supplementary material attached to the analysis, for the signal regions listed in Table I. For the LHCb analysis we assume $\epsilon_{\text{trig}}(v) \approx \epsilon_{\text{off}}(v) \approx 1$, in line with the study in [57].

For searches requiring just a single vertex, we can simply define the weight of the whole event $w(e)$ as

$$w(e) \equiv 1 - \prod_{v \in \{\text{vertices}\}} (1 - w(v)) \quad (7)$$

$$\approx \sum_{v \in \{\text{vertices}\}} w(v), \quad (8)$$

where the approximation is justified whenever $w(v) \ll 1$ for all vertices. For the CMS analysis, we will also consider a signal region where at least two vertices are reconstructed in the event. In this case, the event weight is defined as

$$w(e) \equiv 1 - \prod_{v \in \{\text{vertices}\}} (1 - w(v)) - \sum_{v \in \{\text{vertices}\}} w(v) \prod_{v' \in \{\text{vertices}\} \setminus \{v\}} (1 - w(v')). \quad (9)$$

This reweighting strategy allows us to efficiently compute limits for arbitrary values of $c\tau$, without the need to regenerate the Monte Carlo samples.

The procedure outlined above gives us the signal efficiency as a function of the model parameters. We show the signal efficiency for select benchmark points in Fig. 12 of Appendix B. We can directly compare this with the number of excluded events in the case of CMS [9] or with the limit on the fiducial cross section for LHCb [7] and extract a bound on the branching ratio of the SM Higgs to hidden valley model under consideration for each signal region. CMS and LHCb present their limits binned in terms of L_{xy} and p_T respectively (see Table I). In both cases we take the bin with the strongest limit to represent our limit. CMS only reports results for dimuon pairs whose reconstructed momentum satisfies the pointing cut in Table I. For LHCb, we use the pointing selection for Scenario A and the nonpointing selection for Scenarios B1 and B2.

IV. RESULTS

The relative advantage of a scouting trigger is best seen by looking at the p_T spectrum of the softer of the two muons in each vertex, as shown in Fig. 5 for an example




FIG. 5. The p_T spectra of the trailing muon in each vertex for an example Scenario A model point, within pseudorapidity ranges covered by the CMS tracker and LHCb VELO detectors (red and purple respectively). Dashed lines indicate the p_T threshold of the corresponding analyses, see Sec. III A for details. The dot-dashed lines indicate the threshold for the prompt and displaced dimuon triggers used in respectively [58] and [59,60]. The PYTHIA 8 card used to generate the events in this figure is included in Appendix A 4.




FIG. 6. Limits on the branching ratio for the SM Higgs decaying to the hidden valley model described in Sec. II, with $\Lambda = m_\eta = 4m_{\pi_3}$, $\sin \theta = 0.1$ and $g = 0.05$. (See Appendix A 2 for details.) The bounds are shown as a function of the lifetime of the dark photon $c\tau_{A'}$, from our recast of the CMS dimuon scouting analysis [9] and the LHCb low-mass dimuon analysis [7]. The dashed yellow line indicates the limit set by [9] when requiring two dimuon vertices per event, under the assumption that this suffices to suppress the background to negligible levels. (See text for details.)

working point. While signal events may be muon-rich, their muons are generally soft, and, depending on the parent particle lifetime, very often displaced. The p_T thresholds for displaced muon triggers are prohibitively high for this signal, as Fig. 5 demonstrates, while the scouting trigger can record events with trailing muon p_T as low as 3.5 GeV. Thus, the scouting trigger is uniquely suited to capture the regime where signals yield soft, displaced muon pairs.

For muon pairs with a sufficiently small displacement to be picked up by prompt dilepton triggers, a softer trailing muon p_T cut of 10 GeV is more reflective of ATLAS and CMS capabilities (e.g., [59,60]). But even this p_T threshold still yields an acceptance of only a few percent for the model point in Fig. 5. For the low- $c\tau$ /prompt regime, the scouting analysis in [10] is therefore very pertinent. While the scouting trigger has much better signal acceptance, it also collects more background events than the traditional dimuon trigger, in particular for muon pairs produced at low to no displacement. In some models, the signal may contain additional handles such as jet substructure information or hadronic displaced vertices that could help with background discrimination, but which are not retained in the scouting stream. In such cases, a fully offline analysis may prove to be more powerful. For the benchmark scenarios that we consider in this paper, we do not expect such handles to add much discriminating power in the regime where all dark decays are prompt. We thus suspect that even in the limit of prompt decays the scouting analysis in [10] would outperform the standard dimuon trigger, although this statement is model dependent. We leave a quantitative study of scouting for dark showers in the prompt regime to future work.

A. Scenario A

We first discuss Scenario A, which produces resonant dimuon pairs whose reconstructed momenta point back to the beamline. The bounds we obtain from both LHCb and CMS for benchmark dark shower models produced in exotic Higgs decays is shown in Fig. 6 for a few example mass points. The blue and red shaded regions correspond to the bounds set by the CMS analysis with and without imposing isolation, respectively. The isolation efficiency for these signal benchmarks is roughly 40%, which reflects that in this scenario the dark photons are produced in pairs. We see that imposing isolation is somewhat beneficial for relatively low $m_{A'}$, where the SM backgrounds are substantial. For $1 \text{ GeV} \gtrsim m_{A'} \gtrsim 3 \text{ GeV}$ both selections perform comparably. For $m_{A'} \gtrsim 3 \text{ GeV}$, however, the background drops off sharply [9], and the more inclusive selection sets the best limit. The best sensitivity of the LHCb analysis is at somewhat lower $c\tau_{A'}$, as the dark photons in the LHCb acceptance tend to be more boosted than those in the CMS acceptance. Despite its much lower luminosity, the LHCb analysis outperforms the CMS displaced analysis for $m_{A'} \gtrsim 1 \text{ GeV}$ and $c\tau_{A'} \lesssim 1 \text{ mm}$.³

One of the key features of dark shower topologies is that they tend to produce multiple long-lived particles per event. If the proper decay length of these long-lived particles is comparable to or smaller than the size of the detector and if the vertex reconstruction efficiency is sufficiently high, one expects that requiring an additional displaced dimuon vertex

³The CMS prompt scouting search [10] appeared when our manuscript was in its final stage of completion. It likely has some sensitivity in the low $c\tau$ regime; we defer a detailed analysis of its reach for dark shower models to future work.




FIG. 7. Limits on the Higgs branching ratio to the hidden valley model in Scenario A from our recast of Ref. [9], with $\Lambda = m_\eta = 4m_{\pi_3}$, $\sin\theta = 0.1$, and $g = 0.05$. (See Appendix A 2 for details.) The “single vertex” column indicates the single vertex isolated or inclusive signal region, whichever gives the strongest limit for the point in question. The “double vertex” limit makes important and unvalidated assumptions about correlations in the vertex reconstruction efficiencies and the SM background, as described in the text. It should therefore be interpreted as a sensitivity estimate rather than as a robust limit.

would be a powerful handle to reduce the backgrounds. To illustrate this, we derive a tentative bound from the CMS analysis by assuming that (i) the vertex reconstruction efficiencies $\epsilon_{\text{trig}}(v) \times \epsilon_{\text{off}}(v)$ are independent between different vertices and that (ii) requiring an additional vertex suffices to reduce the background to negligible levels for the whole range of $m_{A'}$. If both assumptions are satisfied, one obtains the yellow curve in Fig. 6. We see that demanding an additional vertex may be a very powerful handle at low to moderate $c\tau_{A'}$. At higher values of $c\tau_{A'}$, a larger fraction of the long-lived A' escape the tracker before decaying, so that experimental sensitivity is dominated by the single-dimuon vertex search.

Figure 7 shows how the sensitivity of the CMS scouting search changes as we vary both m_{π_3} and $m_{A'}$ while keeping m_{π_3}/Λ fixed. Varying m_{π_3} has two effects; a larger m_{π_3} reduces the available phase space for dark mesons, yielding a lower vertex multiplicity. The larger mass however provides a somewhat higher boost to the A' and therefore the muons, increasing the likelihood that said vertices will be reconstructed. Varying $m_{A'}$ primarily affects the $A' \rightarrow \mu\mu$ branching ratio, when it comes to signal acceptance; however, $m_{A'}$ also plays a primary role in determining the background, which is larger at lower dimuon invariant

masses. From Fig. 7 we conclude that the dependence on m_{π_3} is fairly mild, while the results are more sensitive to $m_{A'}$. In particular, we see a dip in sensitivity when $m_{A'}$ is close to the mass of the SM ρ meson, since resonant mixing with the SM ρ enhances the $A' \rightarrow \pi^+\pi^-$ partial width and accordingly suppresses the branching ratio of the dark photon into muons. In this figure we again see that a double vertex analysis is particularly powerful at low $c\tau_{A'}$ and loses sensitivity faster at high $c\tau_{A'}$, as expected.

While we have studied a specific benchmark dark shower model, thanks to the inclusive nature of both the LHCb and CMS searches, we can draw several broad conclusions. In particular, we see that the scouting search provides the leading sensitivity to low-mass dark shower signatures when the muons are produced in the decay of a parent particle with lifetime $c\tau \gtrsim \text{mm}$, and for $c\tau \gtrsim \text{cm}$ the sensitivity is likely dominated by final states with a single dimuon vertex. Since the scouting search depends only on the kinematics and multiplicity of dimuon pairs, these properties will hold broadly across a range of dark shower models that produce low-mass, muonphilic final states, including those considered in [35,36,38]. The relative advantage of single- vs double-vertex searches will depend weakly on the mass of the mediator initiating the shower as

well as the dark gauge coupling, both of which impact the multiplicity of hidden hadrons. However, acceptance alone ensures that in the longer-lifetime regime the single-vertex search will come to dominate the sensitivity. We can thus conclude that scouting searches are a uniquely sensitive tool for discovering low-mass hidden valley models, providing leading sensitivity in the well-motivated but not universal cases where such models give rise to displaced dimuon resonances that point back to the primary vertex. A double-vertex search would moreover provide important additional sensitivity at intermediate parent lifetimes.

B. Scenario B

Both CMS and LHCb have selection criteria that require the vector sum of the three-momenta of the muons in a given vertex to point back to the beamline, as this is a good handle to reduce background from fake vertices and material interactions. However, this pointing cut has poor efficiency on signals where a long-lived state decays to a dimuon pair plus one or more other states. Scenario B presents two such example cases, where we assume that the A' decays promptly, but that the dark meson π_3 (Scenario B1) or η (Scenario B2) has a macroscopic lifetime.

We can characterize the degree of pointing through the variable $\Delta\phi(\sum \vec{p}_T(\mu), \vec{x})$, which at truth-level corresponds to the azimuthal angle $\Delta\phi$ between the three-momenta of the A' and its parent meson, as shown in Fig. 8 for two example points in Scenario B2. Though the distribution of $\Delta\phi(\sum \vec{p}_T(\mu), \vec{x})$ is still peaked at small values, it is rather broad and the CMS analysis cut (dot-dashed line) is only $\mathcal{O}(10\%)$ efficient. LHCb on the other hand also reports results for an inclusive selection that does not require pointing. This selection comes roughly with an order of magnitude more background than the selection requiring




FIG. 8. Truth-level distribution of $\Delta\phi(\sum \vec{p}_T(\mu), \vec{x})$ for two example benchmark points in Scenario B2. The dashed line indicates the cut imposed in the CMS scouting analysis [9]. The efficiency of this cut for the red (blue) benchmark is 15% (10%).




FIG. 9. Limits on the branching ratio for the SM Higgs decaying to the hidden valley models described in Sec. II, with $\Lambda = m_\eta = 4m_{\pi_3}$, $\sin\theta = 0.1$, and $g = 0.05$. (See Appendix A 2 for details.) The bounds are shown as a function of the lifetime of the dark photon $c\tau_{A'}$, from our recast of the CMS dimuon scouting analysis [9] and the LHCb low-mass dimuon analysis [7].

pointing, which is acceptable if a large enough increase in signal efficiency can be achieved.

We show the resulting bounds on both Scenarios B1 and B2 in Fig. 9 for two example mass points. As before, blue and red shaded regions correspond to the bounds set by the CMS analysis with and without imposing isolation, respectively. The isolation efficiency for Scenario B1 is comparable to that for Scenario A, roughly 40%, while for Scenario B2 it is nearly 100%.

Since the multiplicity and kinematics of dark photons produced in Scenarios A and B1 are very similar, the differences between the bounds on these two scenarios are mainly attributable to the dimuon pair not pointing to beamline. Concretely, the LHCb limit on Scenario B1 using the inclusive selection is generally weaker than the Scenario A limit, which uses the pointing selection at the same dark photon mass, owing to larger backgrounds. The relative stringency of the limits based on these two selections varies bin-by-bin, reflecting background fluctuations; for the specific mass point shown in Fig. 9, the Scenario B1 and Scenario A limits are comparable. Meanwhile for the CMS scouting search, the loss of acceptance from the cut on $\Delta\phi(\sum \vec{p}_T(\mu), \vec{x})$ weakens the Scenario B1 limit by roughly a factor of five compared to Scenario A.

For Scenario B2, the number of dimuon pairs per event is substantially smaller, and accordingly the sensitivity of the existing analyses is further reduced. In the bulk of parameter space for this scenario, the existing CMS and LHCb searches do not have better sensitivity than that offered by indirect constraints from global fits to Higgs properties, which currently constrain the exotic Higgs branching fraction to be less than 15% [61]. It would be very interesting for the CMS collaboration to attempt a search with relaxed pointing criteria, as the enhanced signal




FIG. 10. Truth-level invariant mass of dimuon pair in Scenario C, for two benchmark model points with cm-scale proper decay lengths.

efficiency may be sufficient to offset the higher backgrounds, particularly for isolated dimuon pairs.

C. Scenario C

Finally, in Scenario C, long-lived η mesons produce a displaced dimuon pair that is nonresonant and hence necessarily nonpointing. The kinematics of the decay produce an endpoint in the dimuon mass at $m_{\mu\mu, \text{max}} = m_{\eta} - m_{\pi_2}$, as shown in Fig. 10. Because of the off shell dark photon and the three-body final state, the minimum possible η lifetime for this case is a strong function of the mass splitting between the η and the π_2 . For splittings of the order $m_{\eta} - m_{\pi_2} \sim \text{few GeV}$, the minimum possible proper decay length is $\mathcal{O}(\text{cm})$, where the exact value depends on the remaining model parameters (see Appendix A 3). Very narrow splittings ($m_{\eta} - m_{\pi_2} \ll \text{GeV}$) are therefore likely to have intractably low acceptance.

As for Scenario B, we are unable to reliably model the backgrounds for CMS in the absence of the pointing cut and therefore do not attempt to make a sensitivity projection. Of course, it is clear this is a much more challenging signal than Scenario B, as the distribution of the signal events over a range of dimuon masses both dilutes the overall statistical significance of the signal as well as complicates the data-driven background estimation process. An interesting open question at this juncture is to compare the sensitivities of analyses based on scouting triggers to those using alternative trigger strategies, which may have access to a smaller number of signal events, but may contain more information that can be used to control backgrounds.⁴ The parked

⁴Alternatively, one may attempt to expand the amount of useful information recording in the scouting analysis, e.g., by making use of machine learning-driven data compression methods [62].

data set [63] is a particularly intriguing possibility in this context. For the more traditional triggers, the relative sensitivities will necessarily depend in detail on the production mechanism.

As a concrete example, our benchmark scenario of production in exotic Higgs decays offers the chance to trigger on Higgs production in association with a leptonically decaying W or Z boson using prompt lepton triggers.⁵ While the Higgs production cross section in semileptonic processes is two orders of magnitude smaller than the inclusive Higgs production cross section relevant for scouting, events that arrive on a standard lepton trigger will contain information about the whole event and will be subject to different backgrounds. Without being able to reliably model the backgrounds for nonpointing dimuon pairs, we cannot determine which analysis strategy offers the best sensitivity.

We can, however, conclude that the Higgs production channel considered here is perhaps the most optimistic choice when it comes to alternative trigger pathways for low-mass dark showers. Other natural choices for light ($\lesssim \text{few hundred GeV}$) mediators are *i*) the Z boson [36,38] and *ii*) a new SM-singlet, whether vector or scalar (e.g., [35]); one could straightforwardly extend our benchmark model to cover either case by adding additional heavy states, without substantially changing the dark meson phenomenology. For a Z boson mediator, associated WZ production is more suppressed compared to inclusive Z production than VH is relative to inclusive H production; meanwhile for a BSM mediator such as a Z' or a new scalar, the best additional handle that one could generically expect is energetic initial-state radiation (ISR). In this case, the presence of a moderately hard ISR jet could enable a MET or a standard dimuon trigger strategy, depending on the specifics of the model, at the cost of substantially (and perhaps prohibitively) reducing the available signal cross section. An example of this strategy was very recently deployed in the context of a search for inelastic dark matter [65], which relied on a MET trigger and additionally required a soft displaced dimuon pair in the final state. As this search indicates, the thresholds necessary for a MET trigger strategy are not small; Ref. [65] placed an offline analysis cut of $E_T > 200 \text{ GeV}$.

V. SUMMARY AND RECOMMENDATIONS

The search for hidden valleys at the LHC is a challenging and multifaceted program, thanks in part to the

⁵The associated $t\bar{t}H$ production cross section, which is 0.5 pb at $\sqrt{s} = 13 \text{ TeV}$ [64], can provide a sizeable contribution to Higgs production in association with one or more prompt leptons. This production channel is often a challenge for SM analyses owing to the combinatoric challenges of its final states, but it can be important for recording beyond-the-SM Higgs decays to exotic final states.

fantastic diversity of possible signatures. Fortunately, search strategies for models with a dark confinement scale at the few GeV scale admit a certain amount of streamlining; many of the portal operators that could govern decays of GeV-scale dark hadrons to SM final states predict lifetimes that are too long to leave much visible energy within the main detectors [33]. For such low-mass dark shower events, the operators that are easily compatible with high-multiplicity visible final states tend to be muonphilic. Thus, low- p_T (displaced) dimuon pairs stand out as one of the most promising signatures of low-mass hidden valley theories.

Online analysis techniques, including CMS' data scouting stream and the LHCb low mass dimuon searches, are a powerful tools for probing the soft dimuons characteristic of low-mass dark shower events. We establish the sensitivity of CMS' scouting search for displaced dimuon resonances [9] to dark showers produced in exotic decays of the SM Higgs boson, and compare its reach to that of an inclusive dimuon resonance search at LHCb [7]. We demonstrate the significant advancement in reach realized in the scouting search in the context of a benchmark hidden valley model; however, the inclusive nature of the experimental searches lets us make several observations that hold across a broader class of low-mass dark shower signatures. Concretely, CMS performs best for medium to long proper lifetimes due to its larger integrated luminosity, while LHCb currently sets the strongest bounds at short lifetimes. A reinterpretation of the prompt scouting analysis by CMS [10] in terms of hidden valley models would be well-motivated. Though we cannot make sharp statements without more information on the background, we expect that double vertex analysis selection by CMS and potentially by LHCb could further improve the bounds.

Our simple and flexible benchmark model features two flavors of light dark quarks and an elementary dark photon, and is capable of realizing several different dimuon vertex topologies with relatively few parameters. The simplest scenario occurs when dimuon resonances are produced from the decay of a promptly-produced particle, and thus the reconstructed dimuon momentum points back to the beamline. Our model also realizes scenarios where the dimuon pair does not point back to the beamline, as well as a nonresonant displaced dimuon signature from a three-body decay.

For the nonpointing case, the LHCb search currently has a suitable selection, which has only marginal sensitivity with the current dataset. It would therefore be important for both LHCb and CMS to (continue to) include a nonpointing signal region in future searches with more integrated luminosity. Finally, the nonresonant, nonpointing case is also well-motivated and always implies a macroscopic lifetime. We recommend that it too be included as a possibility in future searches. It remains an interesting open question whether scouting searches or searches for (e.g.)

MET in addition to nonpointing dimuon vertices in events that arrive on higher-threshold trigger streams can provide the best probe of these nonresonant scenarios. However here scouting has the particular advantage that searches are independent of assumptions about the production mode for these showers; alternate trigger strategies depend more strongly on the model-dependent features of the mediator responsible for initiating the shower.

Online analyses such as scouting therefore offer an unmatched discovery tool for low-mass dark showers, which can otherwise easily evade traditional detection strategies. Of course, in the event of a compelling excess, recording more features of the events would be a top priority. The relatively high-multiplicity final states that can be realized in our benchmark models can also provide other handles to further identify and characterize the signature, in particular displaced hadronic and/or electronic vertices and a moderate amount of missing energy. Displaced hadronic final states in particular are a locus of discussion for future trigger capabilities at LHCb as well as ATLAS and CMS [66], and could offer further windows onto this challenging class of signatures.

ACKNOWLEDGMENTS

We gratefully acknowledge useful conversations with Hsin-Chia Cheng, Matthew Citron, Zeynep Demiralgi, Lingfeng Li, Zoltan Ligeti, Steven Lowette, Mario Masciovecchio, Titus Mombächer, Dean Robinson, Ennio Salvioni, Christiane Scherb, Matthew Strassler, Indara Suarez, Mahiko Suzuki, Xabier Cid Vidal, and Michael Williams. We thank Matthew Citron, Lingfeng Li, and Ennio Salvioni for comments on the manuscript. S. K. would also like to thank Diego Redigolo, Alberto Mariotti, and Sam Junius for collaboration on related work. The work of S. B. was supported by a grant from the United States-Israel Binational Science Foundation (BSF). The work of S. K. was supported by the U.S. Department of Energy, Office of Science under Contract No. DE-AC02-05CH11231. Part of this work was performed at the Aspen Center for Physics, which is supported by National Science Foundation Grant No. PHY-1607611. J. S. thanks the MIT Center for Theoretical Physics and Laboratory of Nuclear Science for their generous hospitality during the completion of this work.

APPENDIX A: A TWO-FLAVOR HIDDEN VALLEY

Here we construct the benchmark two-flavor model with a light dark photon used in this work. We first discuss the UV Lagrangian in terms of elementary dark sector quarks and dark gluons, then consider the resulting meson spectrum and use chiral perturbation theory to calculate the relevant branching ratios and decay widths. We end with a discussion of the PYTHIA implementation.

1. Quark sector

We consider two dark sector quarks (q_i) and their corresponding antiquarks (\bar{q}_i), which are respectively in the fundamental and the antifundamental representation of the dark sector $SU(N_c)$ strong force. We furthermore introduce a single, weakly coupled $U(1)$ gauge field A' . We will be interested in the case when this gauge field has a mass. For Scenarios A and B1, where the anomalous decay $\pi_3 \rightarrow A'A'$ is responsible for the visible signal, one can take this mass to come from a Stückelberg mechanism, which would be compatible with extending the model into a “neutral naturalness” solution to the hierarchy problem. In order to realize the distinct vertex topologies in Scenarios B2 and C, we additionally include a dark Higgs scalar ϕ , whose vacuum expectation value will contribute to the mass of A' as well as those of the dark quarks. The charge assignments are as follows:

	q_1	q_2	\bar{q}_1	\bar{q}_2	ϕ	
$SU(N_c)$	\square	\square	\square	\square	1	(A1)
$U(1)$	1	-1	0	0	1	

which, at the renormalizeable level, allow for the following set of interactions

$$\mathcal{L} \supset i \sum_i q_i^\dagger \not{D} q_i + i \sum_i \bar{q}_i^\dagger \not{D} \bar{q}_i + |D^\mu \phi|^2 \quad (\text{A2})$$

$$- \begin{pmatrix} q_1 \\ q_2 \end{pmatrix}^T \begin{pmatrix} y_{11} \phi^\dagger & y_{12} \phi^\dagger \\ y_{21} \phi & y_{22} \phi \end{pmatrix} \begin{pmatrix} \bar{q}_1 \\ \bar{q}_2 \end{pmatrix} + \text{H.c.} \quad (\text{A3})$$

$$- \epsilon e A'_\mu J_{\text{EM}}^\mu \quad (\text{A4})$$

with $D_\mu \equiv \partial_\mu + igQA'_\mu + ig_s T^a G_\mu^a$. Here g and g_s are respectively the $U(1)$ and $SU(N_c)$ gauge couplings, with Q and T^a the $U(1)$ charge and the $SU(N_c)$ generators respectively. The G_μ^a are the $SU(N_c)$ gluons. For the mass range of interest here, the dark photon can to good approximation be taken to have a small coupling ϵe to the SM electromagnetic current J_{EM}^μ , through the mixing of the SM photon with the A' . We further assume that the scalar field develops a vacuum expectation value v , such that we can decompose it as

$$\phi = \frac{1}{\sqrt{2}}(v + \varphi)e^{i\tilde{a}/v}. \quad (\text{A5})$$

The vacuum expectation value contributes both to the mass of the A' and the masses of the quarks through the interactions in Eqs. (A2) and (A3). Concretely, expanding the kinetic term for ϕ we find

$$|D^\mu \phi|^2 \supset \frac{1}{2} g^2 v^2 A'_\mu A'^\mu + g v A'_\mu \partial^\mu \tilde{a} \left(1 + \frac{\varphi}{v}\right)^2. \quad (\text{A6})$$

A Higgs portal coupling between SM and dark Higgses is allowed by all symmetries and provides a natural UV completion of the coupling $h q_i \bar{q}_i$ responsible for initiating the dark showers, as we now sketch. Introducing a Higgs portal mixing between dark and SM Higgs of the form $\mathcal{L}_{\text{int}} = \kappa |\phi|^2 |H|^2$ gives rise to a dark Higgs-SM Higgs mixing angle that in the limit $\kappa \ll 1$ can be expressed in terms of the dark Higgs vev v and mass m_φ as

$$\theta \approx \frac{\kappa v v_h}{m_\varphi^2 - m_h^2}, \quad (\text{A7})$$

where v_h is the SM Higgs vev. As a result of this mixing, the SM Higgs h picks up a coupling to dark quarks given by

$$\mathcal{L}_{\text{int}} = \frac{y_{ij}}{\sqrt{2}} \theta h \bar{q}_{Di} q_{Dj}. \quad (\text{A8})$$

We will assume for simplicity that the radial mode φ decouples from the subsequent phenomenology, and neglect it hereafter.⁶ In the absence of chiral symmetry breaking, we could simply remove the Goldstone boson \tilde{a} by working in the unitary gauge. The mass of A' however receives a contribution both from v and from the confining $SU(N_c)$ dynamics, which means that a linear combination of the elementary Goldstone mode \tilde{a} with one of the meson modes will furnish the longitudinal component of the A' in the unitary gauge. This is best treated in chiral perturbation theory, with the meson degrees of freedom; we describe it in the next section.

The quark sector in Eq. (A3) *a priori* leaves us with a lot of freedom and is therefore rather unwieldy. Rather than mapping out the phenomenology of the fully general case, we are interested in picking an example that is simple to parametrize and that generates the signatures we are interested in. For this reason we will unapologetically assume the following relations between the Yukawa couplings in Eq. (A3),

$$y_{11} = y_{22} \quad \text{and} \quad y_{12} = y_{21}. \quad (\text{A9})$$

To track the gauge invariance in the low-energy chiral Lagrangian, it will be useful to retain the explicit dependence on the Goldstone mode \tilde{a} coming from the fundamental Higgs boson. Thus, we write the mass matrix as

$$M = \Phi M_0, \quad (\text{A10})$$

⁶Adding the Higgs portal coupling κ also shifts the dark quark masses by a quantity of order $\kappa v_h^2/m_\varphi^2$. For our purposes we can absorb this shift into a redefinition of the quark mass parameters.

where M_0 is the mass matrix with $\tilde{a} = 0$, and $\Phi \equiv \text{diag}(e^{-i\tilde{a}/v}, e^{i\tilde{a}/v})$ in the gauge basis. If we move to the basis that diagonalizes M_0 , the full mass matrix M then reads

$$\mathcal{L} \supset - \begin{pmatrix} \bar{q}_1 \\ \bar{q}_2 \end{pmatrix}^T \begin{pmatrix} m_1 \cos(\frac{\tilde{a}}{v}) & -im_2 \sin(\frac{\tilde{a}}{v}) \\ -im_1 \sin(\frac{\tilde{a}}{v}) & m_2 \cos(\frac{\tilde{a}}{v}) \end{pmatrix} \begin{pmatrix} q_1 \\ q_2 \end{pmatrix} + \text{H.c.} \quad (\text{A11})$$

with $m_1 \equiv \frac{v}{\sqrt{2}}(y_{11} - y_{12})$ and $m_2 \equiv \frac{v}{\sqrt{2}}(y_{12} + y_{11})$. We explicitly retain the phase \tilde{a}/v , as it will play a role in the chiral perturbation theory calculations in the next section.

At this point it is convenient to switch to Dirac notation, which will make the approximate $SU(2) \times SU(2)$ flavor symmetry manifest and facilitate the matching onto chiral perturbation theory. The model is then specified by (with $\tilde{a} = 0$)

$$\mathcal{L} \supset i \sum_{i=1,2} \bar{Q}_i \not{D} Q_i - m_i \bar{Q}_i Q_i \quad (\text{A12})$$

with the Dirac fermions $Q_i = (q_i \bar{q}_i^\dagger)^T$. In this basis, the interactions of the A' are now manifestly chiral, and the covariant derivative is defined as

$$D_\mu \equiv \partial_\mu + ig Q_L P_L A'_\mu + ig Q_R P_R A'_\mu + ig_s T^a G_\mu^a \quad (\text{A13})$$

with the $P_{L,R}$ the chiral projection operators. The charge matrices $Q_{L,R}$ are obtained by rotating the charges in Eq. (A1) to the basis that diagonalizes M_0 ,

$$Q_L = \begin{pmatrix} & 1 \\ 1 & \end{pmatrix} \quad \text{and} \quad Q_R = 0. \quad (\text{A14})$$

In general, the charge matrices in the M_0 basis can be nonsparse and contain noninteger numbers, as does the CKM matrix in the SM. In other words, the simple form of Eq. (A14) is nongeneric and a consequence of our choices in Eq. (A9). This choice is intended to streamline the analysis in the next section, without qualitatively changing the phenomenology. Note that the A' interaction breaks both C and P symmetry, which is essential to realize the $\eta \rightarrow A' \pi_2$ decay channel. It preserves CP , however, and we will assume that the θ -angle associated with the $SU(N_c)$ dynamics is small enough to not meaningfully affect the phenomenology, as is the case in the SM.

2. Meson sector

We assume that gauge coupling and masses in Sec. A 1 are a small perturbation on the strong dynamics in the dark sector. In this case, the lightest dark sector mesons are well-described as the three pseudo-Goldstone bosons associated

by the breaking of the $U(2) \times U(2) \rightarrow U(2)$ symmetry by the dynamics of the confining gauge group. The mesons ($\tilde{\pi}_i$) therefore make up the adjoint representation of the unbroken $U(2)$, such that we can define the following matrix of meson fields

$$\pi \equiv \sum_{i=0}^3 \lambda_i \tilde{\pi}_i \quad (\text{A15})$$

with the λ_i the Pauli matrices for $i = 1, 2, 3$ and $\lambda_0 = \mathbb{1}_{2 \times 2}$. In other words, the $\tilde{\pi}_i$ refer to the fields in the symmetry eigenbasis, as defined by Eq. (A15). We will reserve the π_i notation for the fields in the mass eigenbasis, as introduced below. We further define the matrix of fields

$$\Sigma \equiv \frac{1}{2} f e^{i\pi/f}. \quad (\text{A16})$$

This field transforms as $\Sigma \rightarrow L^\dagger \Sigma R$ under the $U(2) \times U(2)$ flavor symmetry of the UV theory and as $\Sigma \rightarrow V^\dagger \Sigma V$ under the unbroken $U(2)$. The meson effective theory is then described by

$$\mathcal{L} = \mathcal{L}_{\text{kin}} + \mathcal{L}_{\text{mass}} + \mathcal{L}_{\text{an}} \quad (\text{A17})$$

with

$$\mathcal{L}_{\text{kin}} = \text{Tr}[D^\mu \Sigma^\dagger D_\mu \Sigma] \quad (\text{A18})$$

$$\mathcal{L}_{\text{mass}} = c_m \Lambda f \text{Tr}[M \Sigma] + \text{H.c.} \quad (\text{A19})$$

$$\mathcal{L}_{\text{an}} = \frac{m_0^2}{8} \text{Tr}[\ln \Sigma - \ln \Sigma^\dagger] = -\frac{1}{2} m_0^2 \tilde{\pi}_0^2 \quad (\text{A20})$$

with

$$D_\mu \Sigma \equiv \partial_\mu \Sigma + ig A'_\mu (Q_L \Sigma - \Sigma Q_R). \quad (\text{A21})$$

Here M is the quark mass matrix defined in Eqs. (A10) and (A11). The parameter Λ in Eq. (A19) represents the dark sector's confinement scale, while c_m is an $\mathcal{O}(1)$, dimensionless matching coefficient. We set $c_m = 1$ going forward.⁷ Eq. (A20) reflects the fact the isospin singlet π_0 is not a true Goldstone boson and receives a mass correction due to instanton effects [52], as does the η' in the SM. The coefficient m_0^2 scales as $1/N_c$ in the large N_c limit. We will assume that $m_0 \gg m_{1,2}$ in our analysis. Here we have kept only the leading terms in the momentum and $1/N_c$

⁷We will always present our results in terms of the physical meson masses; as such the value of c_m will only enter when converting the m_{π_i} to the quark masses. The value of c_m therefore does not affect the phenomenology, except when we verify that the Yukawa couplings in Eq. (A3) are perturbative for the meson spectra of our choice.

expansions, an approximation that reproduces the meson spectrum in the SM to within $\sim 20\%$ accuracy (see e.g., [67]).

We can obtain the mass matrix for the mesons by expanding Eqs. (A19) and (A20) to leading nontrivial order in $1/f$ and Eq. (A11) to second order in $1/v$. This results in

$$\mathcal{L} \supset -\frac{1}{2} \bar{m} \Lambda \begin{pmatrix} \tilde{a} \\ \tilde{\pi}_1 \end{pmatrix}^T \begin{pmatrix} \frac{f^2}{v^2} & \frac{f}{v} \\ \frac{f}{v} & 1 \end{pmatrix} \begin{pmatrix} \tilde{a} \\ \tilde{\pi}_1 \end{pmatrix} - \frac{1}{2} \bar{m} \Lambda \tilde{\pi}_2^2 - \frac{1}{2} \begin{pmatrix} \tilde{\pi}_0 \\ \tilde{\pi}_3 \end{pmatrix}^T \begin{pmatrix} m_0^2 + \bar{m} \Lambda & \delta m \Lambda \\ \delta m \Lambda & \bar{m} \Lambda \end{pmatrix} \begin{pmatrix} \tilde{\pi}_0 \\ \tilde{\pi}_3 \end{pmatrix} \quad (\text{A22})$$

with $\bar{m} \equiv m_1 + m_2$ and $\delta m \equiv m_2 - m_1$. The isospin-conserving limit is retrieved by sending $\delta m \rightarrow 0$ and $v \rightarrow \infty$, in which case we find a light, degenerate isospin triplet, analogous to the SM pions, and a heavy isospin singlet, analogous to the SM η . We will however work with broken isospin, i.e., with $\delta m \neq 0$, such that the $\tilde{\pi}_0$ and $\tilde{\pi}_3$ eigenstates mix. We can define the fields in the mass basis as

$$a \equiv \cos \phi \tilde{a} + \sin \phi \tilde{\pi}_1, \quad (\text{A23})$$

$$\pi_1 \equiv \sin \phi \tilde{a} - \cos \phi \tilde{\pi}_1, \quad (\text{A24})$$

$$\pi_2 \equiv \tilde{\pi}_2, \quad (\text{A25})$$

$$\pi_3 \equiv \cos \theta \tilde{\pi}_3 + \sin \theta \tilde{\pi}_0, \quad (\text{A26})$$

$$\eta \equiv \sin \theta \tilde{\pi}_3 - \cos \theta \tilde{\pi}_0, \quad (\text{A27})$$

with the mixing angles specified by

$$\tan \phi = \frac{f}{v}, \quad (\text{A28})$$

$$\tan \theta = \frac{\sqrt{m_0^4 + 4\delta m^2 \Lambda^2} - m_0^2}{2\delta m \Lambda} \quad (\text{A29})$$

$$\approx \frac{\Lambda \delta m}{m_0^2} + \mathcal{O}\left(\frac{\delta m^3 \Lambda^3}{m_0^6}\right). \quad (\text{A30})$$

The masses are

$$m_a^2 = 0, \quad (\text{A31})$$

$$m_{\pi_1}^2 = \bar{m} \Lambda \left(1 + \frac{f^2}{v^2}\right), \quad (\text{A32})$$

$$m_{\pi_2}^2 = \bar{m} \Lambda, \quad (\text{A33})$$

$$m_{\pi_3}^2 = \bar{m} \Lambda \left(1 - \tan \theta \frac{\delta m}{\bar{m}}\right), \quad (\text{A34})$$

$$m_\eta^2 = m_0^2 + \bar{m} \Lambda \left(1 + \tan \theta \frac{\delta m}{\bar{m}}\right). \quad (\text{A35})$$

The massless linear combination a will furnish the longitudinal component of the A' in the unitary gauge, as we show below. The isospin-preserving limit corresponds to $\sin \theta \rightarrow 0$ and $v \rightarrow \infty$.

Turning now to the kinetic term in Eq. (A18), we can add the kinetic term for the Higgs field in Eq. (A6), expand in $1/f$ and subsequently move to the mass eigenbasis. Aside from canonical kinetic terms for the a , π_i , and η fields, this yields the following terms:

$$\mathcal{L} \supset \frac{1}{2} m_{A'}^2 A'^\mu A'_\mu + m_{A'} A'^\mu \partial_\mu a \quad (\text{A36})$$

$$+ g \sin \theta A'^\mu (\pi_2 \partial_\mu \eta - \eta \partial_\mu \pi_2) \quad (\text{A37})$$

$$+ g \cos \theta A'^\mu (\pi_3 \partial_\mu \pi_2 - \pi_2 \partial_\mu \pi_3) + \dots \quad (\text{A38})$$

with $m_{A'} \equiv g\sqrt{v^2 + f^2}$. The \dots represent higher order terms in the $1/v$ and $1/f$ expansions. We see that the A' mass term receives a contribution from the chiral condensate, as expected. The massless linear combination a moreover indeed corresponds to the mode that is eaten by the A' in the unitary gauge.

As long as isospin is broken ($\sin \theta \neq 0$), the $\eta \rightarrow A' \pi_2$ decay mode is available. This decay is both C and P violating, similar to the $\eta' \rightarrow \rho \pi_0$ mode in the SM. Unlike the SM, the UV completion we chose in Appendix A 1 maximally breaks both C and P , such that the $\eta \rightarrow A' \pi_2$ decay in our dark sector is allowed. The decay will always proceed to the longitudinal component of the A' , since the decays to the transverse polarizations are incompatible with angular momentum conservation.

3. Decay rates and branching ratios

We consider three decay modes for the dark sector mesons π_3 and η that can give rise to displaced muon pairs:

- (i) $\pi_3, \eta \rightarrow A' A'$ through the A' chiral anomaly;
- (ii) $\eta \rightarrow \pi_2 A'$ through Eq. (A37);
- (iii) $\eta \rightarrow \pi_2 A'^* \rightarrow \pi_2 f \bar{f}$ through Eq. (A37) as a three-body decay, with the A' off shell.

For simplicity, we will always choose our mass benchmark points such that other pion decay modes such as $\pi_2 \rightarrow \pi_3 A'$ are kinematically closed. The other dark mesons $\pi_{1,2}$ are detector-stable and contribute to missing energy.

The $\eta \rightarrow 3\pi$ decay is isospin-violating, but otherwise proceeds through the dark sector's strong interaction. Provided it is kinematically allowed and $\sin \theta$ is not tiny, we therefore assume that this decay happens promptly. The leading terms in the chiral Lagrangian, Eqs. (A17)–(A20), predict equal branching ratios to the final states $\pi_1 \pi_1 \pi_3$ and $\pi_2 \pi_2 \pi_3$, up to phase space corrections due to the small mass

difference between the pions. However, these branching ratios will be corrected by the higher-order operator $\text{Tr}(D_\mu \Sigma^\dagger D^\mu \Sigma (M \Sigma + \text{H.c.}))$, which contributes at the same order in isospin-breaking and derivatives, and whose coefficient is *a priori* undetermined. For definiteness and simplicity, when the $\eta \rightarrow 3\pi$ mode is kinematically open, we will assume

$$\text{Br}[\eta \rightarrow \pi_1 \pi_1 \pi_3] = \text{Br}[\eta \rightarrow \pi_2 \pi_2 \pi_3] = \frac{1}{2}, \quad (\text{A39})$$

again up to phase-space corrections.

The $\pi_3, \eta \rightarrow A'A'$ decays are mediated by the anomalous current

$$\partial_\mu J_5^{0\mu} = -\frac{\alpha'}{8\pi} N_c \text{Tr}[\lambda^0 Q_L^2] F'^{\mu\nu} \tilde{F}'_{\mu\nu} \quad (\text{A40})$$

$$= -\frac{\alpha'}{4\pi} N_c F'^{\mu\nu} \tilde{F}'_{\mu\nu}, \quad (\text{A41})$$

where in the first line the trace runs only over flavor indices, and $\alpha' \equiv g^2/4\pi$. We can identify the matrix elements by

$$\langle 0 | J_5^{0\mu} | \pi_3 \rangle = -i \sin \theta f_\pi p^\mu e^{-ix \cdot p}, \quad (\text{A42})$$

$$\langle 0 | J_5^{0\mu} | \eta \rangle = -i \cos \theta f_\eta p^\mu e^{-ix \cdot p}, \quad (\text{A43})$$

with f_π and f_η the π_3 and η decay constants. Consequently, we find

$$\Gamma_{\pi_3 \rightarrow A'A'} = \frac{\alpha'^2}{64\pi^3} N_c^2 \sin^2 \theta \frac{m_{\pi_3}^3}{f^2}, \quad (\text{A44})$$

$$\Gamma_{\eta \rightarrow A'A'} = \frac{\alpha'^2}{64\pi^3} N_c^2 \cos^2 \theta \frac{m_\eta^3}{f^2}, \quad (\text{A45})$$

where for definiteness we set $f_\pi \approx f_\eta \approx f$. For the π_3 meson, this is the only decay mode, which proceeds with a proper lifetime of

$$c\tau \sim 1.7 \times 10^{-7} \text{ cm} \times \left(\frac{0.01}{\alpha'} \right)^2 \times \left(\frac{0.1}{\sin \theta} \right)^2 \times \left(\frac{1 \text{ GeV}}{m_{\pi_3}} \right)^3 \times \left(\frac{f}{0.2 \text{ GeV}} \right)^2. \quad (\text{A46})$$

In other words, by dialing the gauge coupling α' , one can choose the π_3 to be either prompt or very long-lived on collider-relevant timescales.

The η can *a priori* also decay through $\eta \rightarrow \pi_2 A'$, provided that $m_{\pi_2} + m_{A'} < m_\eta$. The partial width of this channel is

$$\Gamma_{\eta \rightarrow \pi_2 A'} = \frac{\alpha' \sin^2 \theta m_\eta^3}{4m_{A'}^2} ((1 - x_\pi^2 - x_A^2)^2 - 4x_\pi^2 x_A^2)^{3/2} \quad (\text{A47})$$

with $x_\pi \equiv m_{\pi_2}/m_\eta$ and $x_A \equiv m_{A'}/m_\eta$. The decay width in Eq. (A47) appears to diverge in the $m_{A'} \rightarrow 0$ limit, which is expected from the Goldstone equivalence theorem. Provided that there is no extreme phase space suppression ($x_\pi \ll 1$ and $x_A \ll 1$) and that the $\eta \rightarrow A'A'$ channel is either subleading or kinematically closed, the η lifetime is

$$c\tau \sim 3 \times 10^{-11} \text{ cm} \times \left(\frac{0.01}{\alpha'} \right) \times \left(\frac{0.1}{\sin \theta} \right)^2 \times \left(\frac{2 \text{ GeV}}{m_\eta} \right)^3 \times \left(\frac{m_{A'}}{0.5 \text{ GeV}} \right)^2. \quad (\text{A48})$$

Also here we see that the decay can be made either prompt or displaced, depending on the values adopted for α' and $\sin \theta$.

Finally, if the $\eta \rightarrow \pi_2 A'$ decay is kinematically closed, the three-body decay $\eta \rightarrow \pi_2 f \bar{f}$ will dominate. The partial width to muons for this process is

$$\frac{d\Gamma_{\eta \rightarrow \pi_2 \mu \mu}}{dx} = \frac{\epsilon^2 \alpha' \alpha \sin^2 \theta}{12\pi} m_\eta \frac{x + 2x_\mu^2}{(x - x_A^2)^2} \sqrt{x - 4x_\mu^2} \times \left(\frac{(x-1)^2 - 2(x+1)x_\pi^2 + x_\pi^4}{x} \right)^{3/2}, \quad (\text{A49})$$

with α the SM fine structure constant, $x_\mu \equiv m_\mu/m_\eta$ and $x = q^2/m_\eta^2$, where q^2 is the invariant mass formed by the muon pair. The partial width and total width are obtained by evaluating

$$\Gamma_{\eta \rightarrow \pi_2 \mu \mu} = \int_{4x_\mu^2}^{(1-x_\pi)^2} dx \frac{d\Gamma_{\eta \rightarrow \pi_2 \mu \mu}}{dx}, \quad (\text{A50})$$

$$\Gamma_{\eta \rightarrow \pi_2 f \bar{f}} = \int_{4x_\mu^2}^{(1-x_\pi)^2} dx \frac{d\Gamma_{\eta \rightarrow \pi_2 \mu \mu}}{dx} R(x m_\eta^2), \quad (\text{A51})$$

with $R(q^2)$ the R -ratio, as extracted from experiment [68]. Figure 11 shows the proper decay length and dimuon branching ratio for a set of benchmark points, which were chosen to roughly minimize $c\tau$ without resorting to unphysical choices for underlying couplings. The photon-dark photon mixing parameter ϵ was chosen to satisfy the existing constraints from direct searches for dark photons [53]. We see that this decay tends to be somewhat displaced, with $c\tau \gtrsim 1 \text{ cm}$. The branching ratio to muons hovers around 25% and is relatively insensitive to m_η and m_{π_2} . For completeness, we include all underlying parameters for an example point in Table II.




FIG. 11. Proper decay length and dimuon branching ratio for the three-body decay in Eq. (A51), for a few benchmark model points as indicated in the figure. For $3m_{\pi_2} < m_\eta$, the $\eta \rightarrow 3\pi$ decays will always dominate.

4. PYTHIA 8 implementation

For Monte Carlo implementation of our model, we use the PYTHIA 8 hidden valley module [48,49]. Our model relies on a recent update [30], specifically the new `separateFlav = onflag`, which allows us to set the masses of the pions individually. While our work was in the final stages of preparation, PYTHIA version 8.309 was released, which contains further updates to the hidden valley module. We did not benchmark versions 8.308 and 8.309 against each other, but we did update our PYTHON tool to make it compatible with version 8.309 by adding the new required flag `HiddenValley:setLambda = on`. For version 8.308 this line should be commented out in the PYTHIA configuration cards. In the remainder of this appendix we specify our full settings and mention some approximations and caveats.

Firstly, we note that the dark π^\pm with PDG codes ± 4900211 are each other's antiparticles and therefore degenerate in mass. In our case, the mass eigenstates are the π_1 and π_2 , which need not be degenerate, as the dark $U(1)$ is broken spontaneously. We therefore always work in the regime where $m_{\pi_1} \approx m_{\pi_2}$ and simply identify those states with the ± 4900211 states in the PYTHIA module.

TABLE II. Example model point for which a relatively short-lived $\eta \rightarrow \pi_2\mu^+\mu^-$ decay is realized.

Meson sector	m_{π_1}	0.403 GeV
	m_{π_2}	0.4 GeV
	m_{π_3}	0.396 GeV
	m_η	1 GeV
	$m_{A'}$	1 GeV
	$\sin \theta$	0.1
	f	0.080 GeV
	g	1
	ϵ	5×10^{-4}
	$c\tau$	2.8 cm
Quark sector	m_1	0.072 GeV
	m_2	0.088 GeV
	Λ	1 GeV
	m_0	1 GeV
	v	0.997 GeV
	y_{11}	0.012
	y_{12}	0.11
	$\text{Br}[\eta \rightarrow \pi_2\mu^+\mu^-]$	0.37

There is substantial uncertainty in the hadronization probabilities of the dark sector vector mesons, as well as its η -meson. For the vector mesons, we choose `HiddenValley:probVector=0.75`, such that the hadronization probability matches the naive expectation from the counting of the number of degrees of freedom. For the η meson, we fix `HiddenValley:probKeepEta1 = 1.0`, which corresponds to assuming that the suppression from taking $m_\eta > m_\pi$ in the Lund string model [69,70] is sufficient to model the η hadronization probability.

We set the masses of the vector mesons and the η meson to be equal to the confinement scale Λ and choose the constituent quark masses (4900101 and 4900102) to be $\Lambda + m_q$, in line with the recommendations in [30]. Self-consistent PYTHIA cards for this model can be generated with our publicly available PYTHON code [51].

APPENDIX B: SIGNAL EFFICIENCY PLOTS

We present signal efficiencies for the CMS scouting analysis as a function of the lifetime of the long-lived particle, for four of the example benchmark points discussed in the main text.

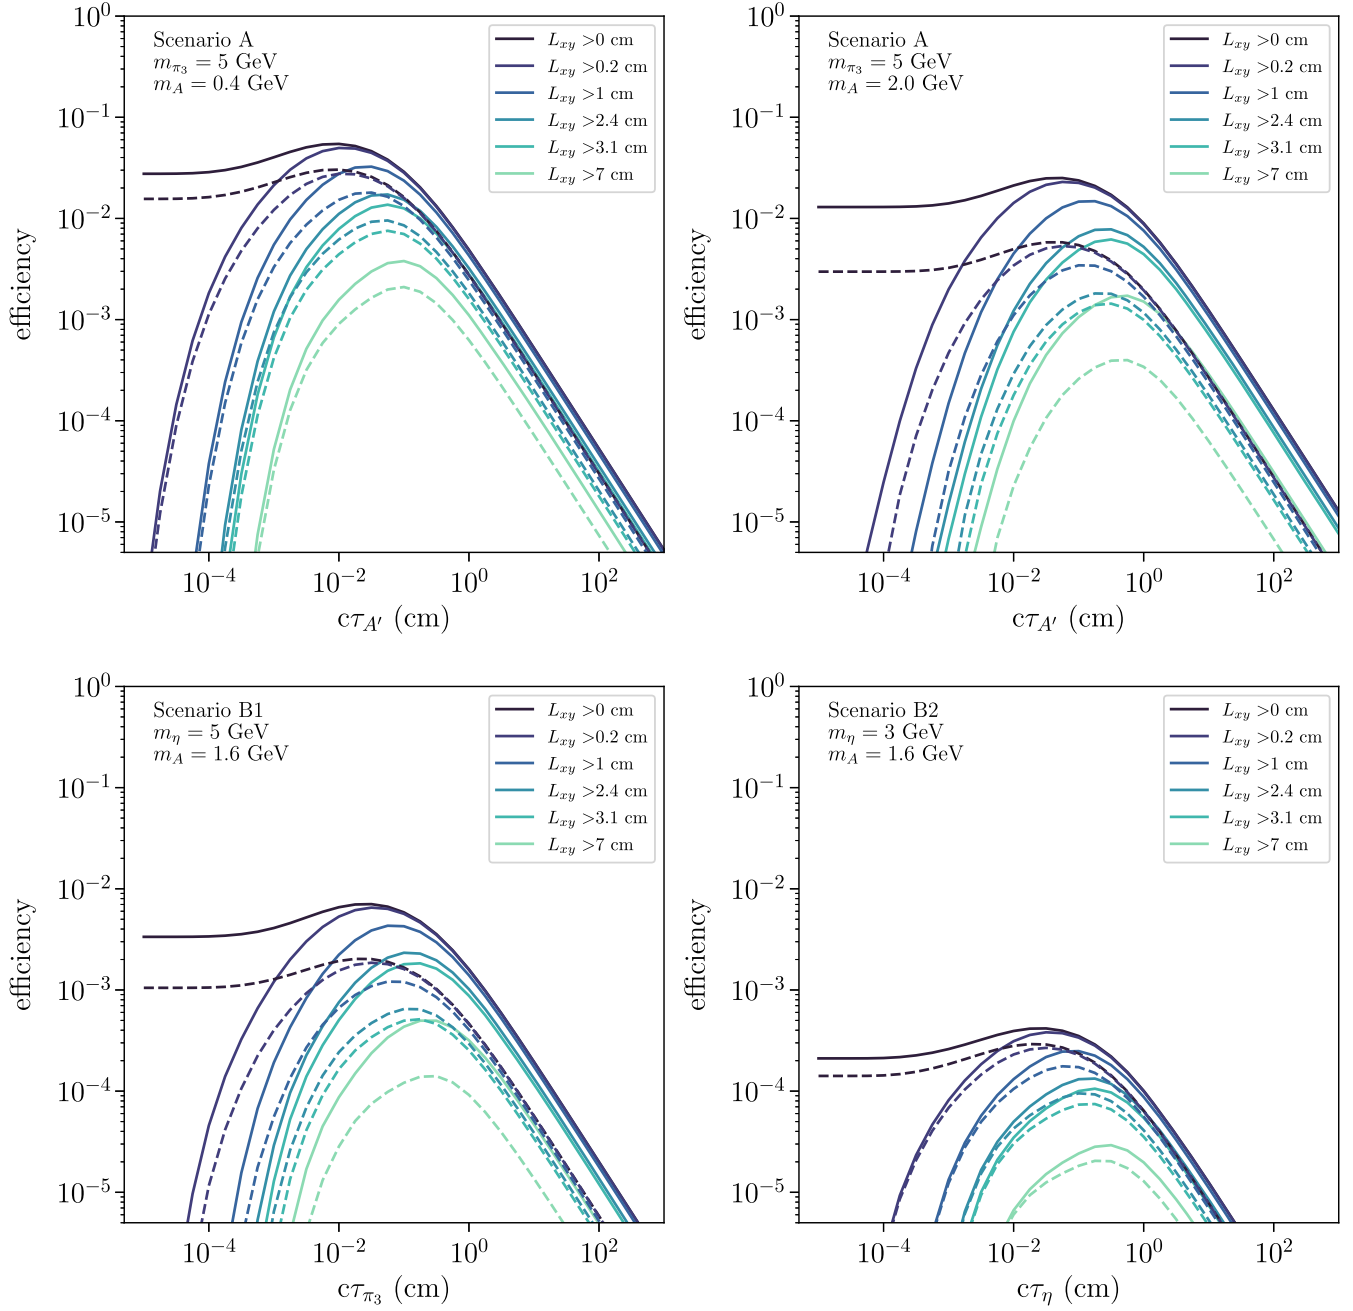


FIG. 12. Signal efficiencies for the CMS analysis, for select benchmark points. Solid (dashed) lines are signal without (with) imposing the isolation condition.

- [1] R. Aaij *et al.*, Tesla: An application for real-time data analysis in high energy physics, *Comput. Phys. Commun.* **208**, 35 (2016).
 [2] R. Aaij *et al.*, A comprehensive real-time analysis model at the LHCb experiment, *J. Instrum.* **14**, P04006 (2019).

- [3] V. Khachatryan *et al.* (CMS Collaboration), Search for Narrow Resonances in Dijet Final States at $\sqrt{s} = 8$ TeV with the Novel CMS Technique of Data Scouting, *Phys. Rev. Lett.* **117**, 031802 (2016).
 [4] M. Aaboud *et al.* (ATLAS Collaboration), Search for Low-Mass Dijet Resonances using Trigger-Level Jets with the

- ATLAS Detector in pp Collisions at $\sqrt{s} = 13$ TeV, *Phys. Rev. Lett.* **121**, 081801 (2018).
- [5] R. Aaij *et al.* (LHCb Collaboration), Search for Dark Photons Produced in 13 TeV pp Collisions, *Phys. Rev. Lett.* **120**, 061801 (2018).
- [6] R. Aaij *et al.* (LHCb Collaboration), Search for $A' \rightarrow \mu^+\mu^-$ Decays, *Phys. Rev. Lett.* **124**, 041801 (2020).
- [7] R. Aaij *et al.* (LHCb Collaboration), Searches for low-mass dimuon resonances, *J. High Energy Phys.* **10** (2020) 156.
- [8] A. M. Sirunyan *et al.* (CMS Collaboration), Search for a Narrow Resonance Lighter than 200 GeV Decaying to a Pair of Muons in Proton-Proton Collisions at $\sqrt{s} = 13$ TeV, *Phys. Rev. Lett.* **124**, 131802 (2020).
- [9] A. Tumasyan *et al.* (CMS Collaboration), Search for long-lived particles decaying into muon pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV collected with a dedicated high-rate data stream, *J. High Energy Phys.* **04** (2022) 062.
- [10] CMS Collaboration, Search for prompt production of a GeV scale resonance decaying to a pair of muons in proton-proton collisions at $\sqrt{s} = 13$ TeV, Technical Report No. CMS-PAS-EXO-21-005, CERN, Geneva, 2023.
- [11] S. Knapen, S. Kumar, and D. Redigolo, Searching for axionlike particles with data scouting at ATLAS and CMS, *Phys. Rev. D* **105**, 115012 (2022).
- [12] J. Alexander *et al.*, Dark sectors 2016 workshop: Community report, [arXiv:1608.08632](https://arxiv.org/abs/1608.08632).
- [13] G. Lanfranchi, M. Pospelov, and P. Schuster, The search for feebly interacting particles, *Annu. Rev. Nucl. Part. Sci.* **71**, 279 (2021).
- [14] M. J. Strassler and K. M. Zurek, Echoes of a hidden valley at hadron colliders, *Phys. Lett. B* **651**, 374 (2007).
- [15] T. Hur, D.-W. Jung, P. Ko, and J. Y. Lee, Electroweak symmetry breaking and cold dark matter from strongly interacting hidden sector, *Phys. Lett. B* **696**, 262 (2011).
- [16] G. D. Kribs, T. S. Roy, J. Terning, and K. M. Zurek, Quirky composite dark matter, *Phys. Rev. D* **81**, 095001 (2010).
- [17] H. Beauchesne, E. Bertuzzo, and G. Grilli Di Cortona, Dark matter in hidden valley models with stable and unstable light dark mesons, *J. High Energy Phys.* **04** (2019) 118.
- [18] A. Francis, R. J. Hudspeth, R. Lewis, and S. Tulin, Dark matter from strong dynamics: The minimal theory of dark baryons, *J. High Energy Phys.* **12** (2018) 118.
- [19] E. Bernreuther, F. Kahlhoefer, M. Krämer, and P. Tunney, Strongly interacting dark sectors in the early Universe and at the LHC through a simplified portal, *J. High Energy Phys.* **01** (2020) 162.
- [20] Y. Bai and P. Schwaller, Scale of dark QCD, *Phys. Rev. D* **89**, 063522 (2014).
- [21] Z. Chacko, H.-S. Goh, and R. Harnik, The Twin Higgs: Natural Electroweak Breaking from Mirror Symmetry, *Phys. Rev. Lett.* **96**, 231802 (2006).
- [22] G. Burdman, Z. Chacko, H.-S. Goh, and R. Harnik, Folded supersymmetry and the LEP paradox, *J. High Energy Phys.* **02** (2007) 009.
- [23] N. Craig, A. Katz, M. Strassler, and R. Sundrum, Naturalness in the dark at the LHC, *J. High Energy Phys.* **07** (2015) 105.
- [24] M. J. Strassler, Why unparticle models with mass gaps are examples of hidden valleys, [arXiv:0801.0629](https://arxiv.org/abs/0801.0629).
- [25] J. Kang and M. A. Luty, Macroscopic strings and 'quirks' at colliders, *J. High Energy Phys.* **11** (2009) 065.
- [26] R. Harnik and T. Wizansky, Signals of new physics in the underlying event, *Phys. Rev. D* **80**, 075015 (2009).
- [27] T. Cohen, M. Lisanti, and H. K. Lou, Semivisible Jets: Dark Matter Undercover at the LHC, *Phys. Rev. Lett.* **115**, 171804 (2015).
- [28] P. Schwaller, D. Stolarski, and A. Weiler, Emerging jets, *J. High Energy Phys.* **05** (2015) 059.
- [29] S. Knapen, S. Pagan Griso, M. Papucci, and D. J. Robinson, Triggering soft bombs at the LHC, *J. High Energy Phys.* **08** (2017) 076.
- [30] G. Albouy *et al.*, Theory, phenomenology, and experimental avenues for dark showers: A snowmass 2021 report, *Eur. Phys. J. C* **82**, 1132 (2022).
- [31] S. Renner and P. Schwaller, A flavoured dark sector, *J. High Energy Phys.* **08** (2018) 052.
- [32] L. Li and Y. Tsai, Detector-size upper bounds on dark hadron lifetime from cosmology, *J. High Energy Phys.* **05** (2019) 072.
- [33] S. Knapen, J. Shelton, and D. Xu, Perturbative benchmark models for a dark shower search program, *Phys. Rev. D* **103**, 115013 (2021).
- [34] W. Bensalem and D. Stolarski, Flavor and CP violation from a QCD-like hidden sector, *J. High Energy Phys.* **02** (2022) 011.
- [35] A. Pierce, B. Shakya, Y. Tsai, and Y. Zhao, Searching for confining hidden valleys at LHCb, ATLAS, and CMS, *Phys. Rev. D* **97**, 095033 (2018).
- [36] H.-C. Cheng, L. Li, E. Salvioni, and C. B. Verhaaren, Light hidden mesons through the Z Portal, *J. High Energy Phys.* **11** (2019) 031.
- [37] M. Zhang, Leptophilic composite asymmetric dark matter and its detection, *Phys. Rev. D* **104**, 055008 (2021).
- [38] H.-C. Cheng, L. Li, and E. Salvioni, A theory of dark pions, *J. High Energy Phys.* **01** (2022) 122.
- [39] C. Cazzaniga and A. de Cosa, Leptons lurking in semi-visible jets at the LHC, *Eur. Phys. J. C* **82**, 793 (2022).
- [40] E. Bernreuther, K. Böse, T. Ferber, C. Hearty, F. Kahlhoefer, A. Morandini, and K. Schmidt-Hoberg, Forecasting dark showers at Belle II, *J. High Energy Phys.* **12** (2022) 005.
- [41] H. M. Lee and M.-S. Seo, Communication with SIMP dark mesons via Z' -portal, *Phys. Lett. B* **748**, 316 (2015).
- [42] Y. Hochberg, E. Kuflik, and H. Murayama, SIMP spectroscopy, *J. High Energy Phys.* **05** (2016) 090.
- [43] A. Berlin, N. Blinov, S. Gori, P. Schuster, and N. Toro, Cosmology and accelerator tests of strongly interacting dark matter, *Phys. Rev. D* **97**, 055033 (2018).
- [44] A. Katz, E. Salvioni, and B. Shakya, Split SIMPs with decays, *J. High Energy Phys.* **10** (2020) 049.
- [45] D. Curtin *et al.*, Exotic decays of the 125 GeV Higgs boson, *Phys. Rev. D* **90**, 075004 (2014).
- [46] M. Cepeda, S. Gori, V. M. Outchoorn, and J. Shelton, Exotic Higgs decays, *Annu. Rev. Nucl. Part. Sci.* **72**, 119 (2022).
- [47] M. J. Strassler and K. M. Zurek, Discovering the Higgs through highly-displaced vertices, *Phys. Lett. B* **661**, 263 (2008).
- [48] L. Carloni and T. Sjöstrand, Visible effects of invisible hidden valley radiation, *J. High Energy Phys.* **09** (2010) 105.

- [49] L. Carloni, J. Rathsman, and T. Sjöstrand, Discerning secluded sector gauge structures, *J. High Energy Phys.* **04** (2011) 091.
- [50] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands, An introduction to PYTHIA 8.2, *Comput. Phys. Commun.* **191**, 159 (2015).
- [51] <https://gitlab.com/simonknapen/dark-showers-tool>.
- [52] E. Witten, Large N chiral dynamics, *Ann. Phys. (N.Y.)* **128**, 363 (1980).
- [53] S. Gori *et al.*, Dark sector physics at high-intensity experiments.
- [54] R. Aaij *et al.* (LHCb Collaboration), Search for hidden-sector bosons in $B^0 \rightarrow K^{*0} \mu^+ \mu^-$ decays, *Phys. Rev. Lett.* **115**, 161802 (2015).
- [55] R. Aaij *et al.* (LHCb Collaboration), Search for long-lived scalar particles in $B^+ \rightarrow K^+ \chi(\mu^+ \mu^-)$ decays, *Phys. Rev.* **D95**, 071101 (2017).
- [56] S. Knapen and S. Lowette, A guide to hunting long-lived particles at the LHC, [arXiv:2212.03883](https://arxiv.org/abs/2212.03883).
- [57] P. Ilten, J. Thaler, M. Williams, and W. Xue, Dark photons from charm mesons at LHCb, *Phys. Rev. D* **92**, 115017 (2015).
- [58] CMS Collaboration, Search for long-lived particles decaying to a pair of muons in proton-proton collisions at $\sqrt{s} = 13$ TeV, *J. High Energy Phys.* **05** (2023) 228.
- [59] A. Tumasyan *et al.* (CMS Collaboration), Search for low-mass dilepton resonances in Higgs boson decays to four-lepton final states in proton-proton collisions at $\sqrt{s} = 13$ TeV, *Eur. Phys. J. C* **82**, 290 (2022).
- [60] ATLAS Collaboration, Search for a new Z' gauge boson in 4μ events with the ATLAS experiment, *J. High Energy Phys.* **07** (2023) 090.
- [61] ATLAS Collaboration, Combined measurements of Higgs boson production and decay using up to 139 fb^{-1} of proton-proton collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS experiment, Technical Report No. ATLAS-CONF-2021-053, Geneva, 2021.
- [62] J. H. Collins, Y. Huang, S. Knapen, B. Nachman, and D. Whiteson, Machine-learning compression for particle physics discoveries, [arXiv:2210.11489](https://arxiv.org/abs/2210.11489).
- [63] CMS Collaboration, Recording and reconstructing 10 billion unbiased b hadron decays in CMS, Tech. Rep. CMS-DP-2019-043, Geneva, 2019.
- [64] D. de Florian *et al.* (LHC Higgs Cross Section Working Group), *Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector* (CERN, Geneva, 2017).
- [65] CMS Collaboration, Search for inelastic dark matter in events with two displaced muons and missing transverse momentum in proton-proton collisions at $\sqrt{s} = 13$ TeV, Tech. Rep. CMS-PAS-EXO-20-010, CERN, Geneva, 2023.
- [66] D. Acosta *et al.*, Review of opportunities for new long-lived particle triggers in run 3 of the large hadron collider, [arXiv:2110.14675](https://arxiv.org/abs/2110.14675).
- [67] C. Degrande and J. M. Gerard, A Theoretical determination of the $\eta - \eta'$ mixing, *J. High Energy Phys.* **05** (2009) 043.
- [68] R. L. Workman *et al.* (Particle Data Group), Review of particle physics, *Prog. Theor. Exp. Phys.* **2022**, 083C01 (2022).
- [69] B. Andersson, G. Gustafson, G. Ingelman, and T. Sjöstrand, Parton fragmentation and string dynamics, *Phys. Rep.* **97**, 31 (1983).
- [70] B. Andersson, G. Gustafson, and B. Soderberg, A general model for jet fragmentation, *Z. Phys. C* **20**, 317 (1983).