

This is the accepted manuscript made available via CHORUS. The article has been published as:

Wilsonian effective field theory of two-dimensional Van Hove singularities

Anton Kapustin, Tristan McKinney, and Ira Z. Rothstein

Phys. Rev. B **98**, 035122 — Published 17 July 2018

DOI: [10.1103/PhysRevB.98.035122](https://doi.org/10.1103/PhysRevB.98.035122)

Wilsonian effective field theory of 2D van Hove singularities

Anton Kapustin,^{1,*} Tristan McKinney,^{1,†} and Ira Z. Rothstein^{2,‡}

¹*Walter Burke Institute for Theoretical Physics, California Institute of Technology, Pasadena, CA 91125*

²*Department of Physics, Carnegie Mellon University, Pittsburgh, PA 15213*

(Dated: May 11, 2018)

We study 2D fermions with a short-range interaction in the presence of a van Hove singularity. It is shown that this system can be consistently described by an effective field theory whose Fermi surface is subdivided into regions as defined by a factorization scale, and that the theory is renormalizable in the sense that all of the counterterms are well defined in the IR limit. The theory has the unusual feature that the renormalization group equation for the coupling has an explicit dependence on the renormalization scale, much as in theories of Wilson lines. In contrast to the case of a round Fermi surface, there are multiple marginal interactions with nontrivial RG flow. The Cooper instability remains strongest in the BCS channel. We also show that the marginal Fermi liquid scenario for the quasiparticle width is a robust consequence of the van Hove singularity. Our results are universal in the sense that they do not depend on the detailed properties of the Fermi surface away from the singularity.

PACS numbers: 71.10.Hf, 11.10.Gh, 71.18.+y, 74.20.-z

I. INTRODUCTION

In the 1990s and early 2000s, extensive theoretical work was devoted to the study of systems of fermions in 2D with the Fermi level close to a van Hove singularity^{1–12}. In such a system, the Fermi velocity vanishes at isolated points on the Fermi surface which we will refer to as van Hove points. From a theoretical standpoint, the van Hove singularity is one of the simplest situations in which deviations from standard Fermi liquid theory are expected. For example, the leading order computation of the self-energy^{1,2} shows that with a short-range interaction, the width of the quasiparticles is linear in the energy, a characteristic behavior of the Marginal Fermi Liquid (MFL)¹³. Since the MFL paradigm has been proposed to explain some peculiar properties of the normal state of high- T_c superconductors, it was speculated that high- T_c superconductors are special due to their proximity to a van Hove singularity^{2,14}. While this scenario has fallen out of favor, understanding the effect of van Hove singularities on the Fermi liquid remains an important problem.

Most of the studies cited above focus on the 2D Hubbard model on a square lattice at or near half-filling because of its relevance to cuprate superconductors. In this model, the Fermi surface is diamond-shaped and features two inequivalent van Hove points (i.e. points where the Fermi velocity vanishes) as well as nesting. These features complicate the analysis, and it is hard to disentangle the effects of van Hove points and nesting. In this paper we study in detail the case of a single van Hove point from the point of view of Effective Field Theory (EFT). When applied to the case of a nonsingular Fermi surface, the EFT approach explains the ubiquity of both the Fermi liquid and BCS-type superconductivity^{15–18}.

As was previously noticed in^{4,6}, the hyperbolic dispersion law characteristic of electrons near a 2D van Hove point leads to additional divergences not regulated by

the Wilsonian cutoff Λ , and necessitates the introduction of an additional regulator which we take to be a Fermi velocity cutoff Υ . Υ also plays the role of a factorization scale which splits the Fermi surface into two regions, $v_F > \Upsilon$ and $v_F < \Upsilon$, where two different power counting schemes apply. In each region momenta are split into large “label” momenta and small “residual” components. Previous work on the 2D van Hove singularity has been plagued by nonlocal divergences, and a recent detailed study¹⁹ concluded that the van Hove EFT is not renormalizable when the Fermi level is exactly at the van Hove singularity and has a very narrow range of applicability when the Fermi level is close to it. However, as we show, when momenta are properly power counted, all of the counterterms are independent of the residual momenta in each respective region and therefore should be considered local. Furthermore, the coupling in each region can only depend upon the label momenta. The coupling can be assumed to be independent of momenta only when all components of the momenta are smaller than Λ/Υ .

In the BCS channel, the RG equation for the coupling function explicitly depends on the logarithm of the ratio of the Wilsonian cutoff Λ to the bandwidth W and leads to the well known double logarithmic running^{4,6,8}. This dependence on the UV scale W represents a form of UV/IR mixing and has interesting consequences discussed below.

The situation is reminiscent of high energy scattering processes, such as the Sudakov form factor, where the phase space of gauge bosons is split into two regions which dominate the IR behavior. This splitting leads to additional (rapidity) divergences which necessitate a new regulator²⁰ to distinguish between soft and collinear modes. Summing contributions from these two sectors leads to a cancellation of the regulator but, as in the present case, the cancellation leaves behind a log in the beta function. This in turn leads to double logs in the renormalization group flow.

We utilize our results to study how a van Hove singularity modifies the low energy behavior. In particular, we discuss the Cooper instability and the range of applicability of the Marginal Fermi Liquid scenario. We show that the Cooper instability is the strongest in the BCS channel, as in the case of the circular Fermi surface, but is also present for other kinematic configurations. This happens because of additional marginal interactions which lead to a breakdown of the Fermi liquid picture. We also show that a certain generalization of the MFL scenario is a robust consequence of the van Hove EFT.

II. A TOY MODEL OF A VAN HOVE SINGULARITY

In the 2D Hubbard model on a square lattice, there are two VH points in the Brillouin zone: $p_{VH} = (0, \pi)$ and $p_{VH} = (\pi, 0)$. When the hopping parameters in the x and y directions are not equal, the energies of these two points are different. If the Fermi level is much closer to one than the other, the effective field theory of a single VH singularity should apply. At both of the VH points, $2p_{VH} \sim 0$. We assume there is a unique VH point in the Brillouin zone and time-reversal (T) symmetry, which takes $p \mapsto -p$, is present. Therefore the singularity sits at the origin, a fixed point under T .

Such a scenario is realized by expanding the nearest-neighbor Hubbard model Hamiltonian around the point $p = 0$ to lowest order in momentum components and assuming a zero-range interaction. The resulting action is

$$S = \int d^3x \left[\psi^\dagger (i\partial_t - \varepsilon(-i\nabla) + \mu) \psi - \frac{g}{2} (\psi^\dagger \psi)^2 \right], \quad (1)$$

where the dispersion relation is

$$\varepsilon(p) = p^2 \equiv t_x p_x^2 - t_y p_y^2 \quad (2)$$

and is unbounded from below. p^2 denotes the square of the 2D vector p with respect to the indefinite metric $\text{diag}(t_x, -t_y)$. It is convenient to set $t_x = t_y = 1$ by rescaling p_x and p_y , such that metric becomes $\text{diag}(1, -1)$, and absorbing a factor of $1/\sqrt{t_x t_y}$ into g . If we regard p_x, p_y as periodic with period of order k_B , then t_x, t_y are of order W/k_B^2 , where W is the bandwidth.

As usual, all states with $\varepsilon(p) < \mu$ are assumed to be occupied, so in the free ($g = 0$) limit the excitations of the system are particles and holes, both with nonnegative energy. When the Fermi level μ vanishes, the system has a discrete symmetry $\psi \leftrightarrow \psi^\dagger, x \leftrightarrow y$ which exchanges particles and holes. Furthermore, the quadratic dispersion relation has $O(1, 1)$ invariance, and the short-range interaction preserves this symmetry. Also, for $\mu = 0$, the action (1) is invariant under dilatations

$$\psi(t, x) \rightarrow \lambda^{-1} \psi(\lambda^2 t, \lambda x). \quad (3)$$

Invariance with respect to Galilean boosts is spontaneously broken by the Fermi sea for all values of μ . As

usual, the dilatation symmetry is anomalous on the quantum level. Internal symmetries include $U(1)$ particle-number symmetry and $SU(2)$ spin symmetry.

The interaction term in (1) has zero range, and in momentum space corresponds to a four-point vertex with no momentum dependence. A naive justification for this simple ansatz is that any vertex with more than four fermionic fields or polynomial momentum dependence is irrelevant in the RG sense. Here we assume that under the RG transformations the momenta scale as

$$p_x \rightarrow \lambda p_x, \quad p_y \rightarrow \lambda p_y, \quad (4)$$

so energy has scaling dimension 2 and ψ has scaling dimension 1. Then the chemical potential μ is relevant, the coupling g is marginal, and more complicated interactions are irrelevant.

This naive argument is, as well known, incorrect, due to the fact that momenta tangent to the Fermi surface should not scale under RG flow. Moreover, the theory defined with a contact interaction, eq. (1), is not consistent, as corrections to the four-point function include nonlocal divergences which cannot be absorbed into a renormalization of g ^{4,6}. These divergences can be traced to the noncompactness of the Fermi surface.

III. SETTING UP THE VAN HOVE EFT

To make the theory (1) well defined, one must impose a cutoff on momenta which will render the Fermi surface compact. This cutoff is imposed in addition to the Wilsonian cutoff $|\varepsilon(\mathbf{p})| \leq \Lambda$. We assume Λ is much smaller than the bandwidth $W \sim k_B^2$. We also assume that $|\mu| \ll \Lambda$, so the modes near the Fermi surface are not integrated out.

Let Υ denote this momentum cutoff. The largest possible value for Υ is of order k_B , the size of the Brillouin zone, and thus it is natural to assume that $\Lambda \ll \Upsilon^2$. The region

$$|p_\pm| \leq \Upsilon, \quad |p_+ p_-| \leq \Lambda, \quad (5)$$

where $p_+ = p_x + p_y$ and $p_- = p_x - p_y$, will be called the VH region.²⁴ Within this region, the dispersion law is

$$\varepsilon = p_+ p_-. \quad (6)$$

We are using Λ and Υ as both explicit regulators and factorization scales. Υ has a natural value of order V_F , the typical value of the Fermi velocity away from the VH points, but it can also be chosen parametrically smaller. In any physical result the dependence on Υ should cancel, since its role is merely to separate the VH and NVH regions. On the other hand, in any physical prediction Λ is a placeholder for the RG scale.

The VH region is the part of the Λ -neighborhood of the Fermi surface that is close to the saddle point. In this region, the dispersion relation (6) implies that if




FIG. 1. An example of the division of the Fermi surface into van Hove and non-van Hove regions.

both components of momentum are of the same order, then $p_{\pm} \sim \sqrt{\Lambda} \ll \Upsilon$. In addition to these “soft modes,” the VH region is populated by collinear and anticollinear modes whose momenta scale as $(\Upsilon, \Lambda/\Upsilon)$ and $(\Lambda/\Upsilon, \Upsilon)$ respectively.

The rest of the Λ -neighborhood of the Fermi surface will be called the NVH region. Within this region, the dispersion law is the standard

$$\varepsilon(p) = p_{\perp} v_F(p_{\parallel}), \quad (7)$$

where p_{\perp}/p_{\parallel} are normal/tangential to the Fermi surface. We assume that the NVH region is “featureless,” in the sense that the Fermi velocity does not change too much there, and that it is free of nesting. The first assumption simply means that there are no other van Hove singularities nearby, while the importance of the second assumption will be discussed in Section VIII. Fig. (1) illustrates the division of a representative Fermi surface into the VH and NVH regions.

In general, loop calculations involving modes from the VH region alone will depend on Υ in such a way that the $\Upsilon \rightarrow \infty$ limit leads to additional divergences. Thus, a sensible EFT must include both the VH region and the NVH region. We use the term “full theory” for such an EFT. We make no assumptions about the shape of the Fermi surface in the NVH region. As we will show below, our results are universal to leading log accuracy in the sense that they only depend upon V_F , the typical Fermi velocity in the NVH region, and not the detailed shape of the Fermi surface. Therefore, our results apply to any system with a VH singularity near the Fermi surface that is weakly coupled at energies of order the bandwidth.

We will denote the fields annihilating electrons in the VH and NVH regions ψ_V and ψ_N respectively. The interaction part of the action is

$$S_{int} = \int dt \prod_{i=1}^4 d^2 p_i \sum_{\alpha\beta\gamma\delta} g_{\alpha\beta\gamma\delta} \psi_{\alpha}^{\dagger} \psi_{\beta}^{\dagger} \psi_{\gamma} \psi_{\delta}, \quad (8)$$

where the indices $\alpha, \beta, \gamma, \delta$ take values V and N . In general, $g_{\alpha\beta\gamma\delta}$ can depend on the momenta p_i of the interacting modes. The couplings must match onto each other as the VH modes approach the NVH region and vice versa. For example, g_{NNVV} must match onto g_{VVVV} as the momenta of the first and second modes approach the VH region.

Naively, in light of the dispersion laws (6) and (7), one might think that the coupling functions in (8) should only depend on the p_{\parallel} of the NVH modes and that the only marginal interaction between the VH modes should be a momentum-independent constant. We will see in the next section that this is not self-consistent: one-loop calculations imply that the couplings must depend on momentum in a more generic manner. This is because when both the rapidity cutoff Υ and the Wilsonian cutoff Λ are present, a low momentum scale Λ/Υ also plays a role. We will call Λ/Υ the ultrasoft scale.

We can achieve some simplification by recalling that momentum and energy conservation limits the interactions of the NVH modes to special kinematic configurations¹⁶. These configurations correspond to forward scattering and back-to-back (BCS) scattering. This implies that interactions between NVH modes and VH modes are of two kinds: (1) forward scattering between a VH mode and an NVH mode and (2) scattering of nearly back-to-back VH modes to nearly back-to-back NVH modes and vice versa. As a result, the numbers of VH and NVH particles are separately conserved.

IV. THE ONE-LOOP BETA FUNCTION

A. Generic kinematic configuration

Consider the scattering of VH modes in a generic kinematic configuration. Conservation of momentum implies the NVH modes will not contribute. Thus tree-level interactions are described by a single coupling function of three independent VH momenta. We would like to determine how this function is renormalized.

It is enlightening to first assume that the coupling is a momentum-independent constant, as naive power counting suggests. The manner in which this assumption fails will show us how to appropriately modify the theory.

We subdivide the VH region into three parts: the soft region, where $p_{\pm} \sim \sqrt{\Lambda}$; the collinear region, where $p_{+} \sim \Upsilon$ and $p_{-} \sim \frac{\Lambda}{\Upsilon}$; and the anticollinear region, where $p_{-} \sim \Upsilon$, and $p_{+} \sim \frac{\Lambda}{\Upsilon}$. Fig. 2 illustrates the location of these subregions. This separation is useful for categorizing the contributions to the beta function. Since in this subsection the kinematic configuration is assumed to be generic, the differences and sums of external momenta are of the same order as the momenta themselves.

As usual, we have three diagrams at one-loop level, which we refer to as s -channel (\mathcal{A}_S), t -channel (\mathcal{A}_T), and u -channel (\mathcal{A}_U); see Fig. 3. These three diagrams depend on $K = p_1 + p_2$, $Q = p_1 - p_3$, and $Q' = p_1 - p_4$




FIG. 2. Subdivision of the VH region.

respectively, and each contributes independently to the beta function. The u -channel diagram is identical to the t -channel diagram if we take $Q \leftrightarrow Q'$, so we focus on the t - and s -channel diagrams.




FIG. 3. The diagrams contributing to the renormalization of the coupling at one loop. Not shown is the u -channel diagram, which is given by interchanging the final state particles in the t -channel diagram.

We find (see Appendix) that any one-loop diagram where a collinear external mode and an anticollinear external mode meet at a vertex leads to a power-suppressed contribution to the beta function. This is because the K or Q involved in the interaction always sets a large energy scale which acts to suppress the associated diagram.

Generic t -channel diagrams that do not involve collinear-anticollinear vertices make order-one contributions to the beta function. For example, for a generic interaction between soft modes,

$$\Lambda \frac{d\mathcal{A}_T}{d\Lambda} = \frac{g^2}{4\pi^2} \quad (9)$$

plus power-suppressed terms. There are exceptions in certain special kinematic configurations; see Section IV B.

The behavior of the s -channel diagrams is more complicated. Defining

$$\varepsilon_K = K_+ K_-, \quad (10)$$

we find that generic the s -channel diagrams that do not involve collinear-anticollinear vertices interpolate between being log enhanced when $\varepsilon_K \ll \Lambda$ and order one when $\varepsilon_K \sim \Lambda$. As an example, for generic interactions between soft modes,

$$\Lambda \frac{d\mathcal{A}_S}{d\Lambda} = -\frac{g^2}{4\pi^2} \log \left(\frac{\Lambda}{\varepsilon_K} \right) \quad (11)$$

plus suppressed terms. To avoid confusion, we note that ε_K is not the net energy of the incoming particles.

B. Special kinematic configurations

Eq. (11) appears to imply that the beta function diverges as ε_K approaches zero, thus necessitating the existence of a nonlocal counterterm, which would mean the formalism lacked a systematic power-counting scheme. However, (11) does not apply in the $\varepsilon_K \rightarrow 0$ limit. The divergent behavior is an unphysical artifact of taking the van Hove region to be infinite in extent. If we take the rapidity cutoff Υ into account, we find that when one component of K , say K_- , satisfies

$$|K_-| < \frac{\Lambda}{\Upsilon}, \quad (12)$$

such as for an interaction between only collinear modes, then

$$\Lambda \frac{d\mathcal{A}_S}{d\Lambda} = -\frac{g^2}{4\pi^2} \log \left(\frac{\Upsilon}{K_+} \right) \quad (13)$$

plus order-one terms. If both components of K are ultrasoft (i.e. smaller in magnitude than Λ/Υ), we find to leading log order

$$\Lambda \frac{d\mathcal{A}_S}{d\Lambda} = -\frac{g^2}{4\pi^2} \log \left(\frac{\Upsilon^2}{\Lambda} \right). \quad (14)$$

We can summarize the detailed behavior of the s -channel contribution to the beta function in the following manner:²⁵

$$\Lambda \frac{d\mathcal{A}_S}{d\Lambda} = \begin{cases} -\frac{g^2}{4\pi^2} \log \left(\frac{\Lambda}{\max(K_+, \frac{\Lambda}{\Upsilon}) \max(K_-, \frac{\Lambda}{\Upsilon})} \right), & \varepsilon_K \lesssim \Lambda \\ \mathcal{O}(1) \times \frac{\Lambda}{\varepsilon_K} g^2, & \varepsilon_K \gtrsim \Lambda. \end{cases} \quad (15)$$

If $K_+ \sim \Upsilon$, the log in (13) will not be large, and hence the order-one “corrections” cannot be ignored. As a result, the dependence on K_+ becomes complicated. Similarly, if one component of Q is large while the other is ultrasoft, the t -channel diagram has a complicated dependence on the large component (though unlike the s -channel diagram, it can never become log enhanced). These cases are discussed in more detail in Section VIII. Finally, the t -channel contribution to the beta function vanishes if both components of Q are ultrasoft.

C. Binning and leading-log behavior

At first glance, the behavior of the beta function implied by the above results is rather odd. The contribution from the s -channel diagram in Eq. (15) sometimes depends nonanalytically on the momentum, and the functional form of the results change when the components of K or Q pass a particular threshold (around the scale Λ/Υ). Previous authors¹⁹ have particularly regarded the behavior of the t -channel diagram as a sign of unavoidable nonlocality in the theory. However, as discussed in the next section, similar behavior appears already for a circular Fermi surface, and is dealt with using bins in momentum space of size Λ/K_F . This notion of binning allows for a clear separation between large and small momenta, and was previously used in the context of the theory of non-relativistic heavy quarks²¹. Binning is also implicit in the standard Fermi-surface RG¹⁶. We apply the same method here.

We divide momentum space into bins of size Λ/Υ , each with a label momentum corresponding to the center of the bin and a residual momentum, of order Λ/Υ , corresponding to the position within the bin. The couplings are then indexed by the discrete label momenta, and we can Taylor expand in the residual momenta. The beta function then depends at leading order on the label momenta alone, and all results are analytic in the residual momenta. The theory is therefore renormalizable, although the couplings depends in an arbitrary way on the label momenta. The same is true for a circular Fermi surface (see the next section).

The non-analytic dependence on the net momentum implies that our assumption of a momentum-independent coupling was inconsistent, and the RG flow will generate dependence on the label momenta even for modes within the soft region. While this complicated behavior threatens the predictive power of the theory, we will see in Section VI that the enhancement of the beta function for modes with small net momentum allows for several important simplifications.

V. REVISITING THE ROUND FERMI SURFACE

Let us revisit some old results involving a round Fermi surface. In that context, previous authors^{15,16,17,18} found that only certain coupling functions are present in the IR theory. In particular, only forward scattering and interactions between back-to-back particles (the BCS channel) are marginal, in the language of effective field theory. Furthermore, these authors found that only the BCS coupling is renormalized (that is, corrections to forward scattering are power suppressed), and that the beta function for the BCS interaction is in fact one-loop exact for generic round Fermi surfaces^{16,18}.

These results hold in the limit where the Wilsonian cutoff Λ on the energy of the modes included in the the-

ory (or, in other words, the “width” of modes around the Fermi surface) is taken to zero while the size of Fermi surface itself is held fixed. For nonzero Λ , near-forward and near-BCS scattering continue to be present in the theory. To understand their role more precisely, let us consider their contributions to the one-loop beta function.

We may parameterize a generic coupling function in terms K , Q , and Q' , the same functions of the external momenta defined in Section (IV). As before, the s -, t -, and u -channel diagrams (¹⁶ calls these the BCS, ZS, and ZS' diagrams), depend on K , Q , and Q' respectively, and the t - and u -channel diagrams are exchanged under $Q \leftrightarrow Q'$. BCS scattering occurs for $K = 0$ and forward scattering occurs when either Q or Q' is zero.

It is straightforward to show that when any of these momenta are order K_F (the radius of the Fermi surface), the presence of the Fermi surface forces the contribution from the corresponding one-loop diagram to the beta function to be suppressed. For example, the log derivative of the one-loop s -channel diagram is

$$\Lambda \frac{d\mathcal{A}_s}{d\Lambda} \sim \frac{\Lambda}{v_F |K|} g^2 \quad (16)$$

when $|K| \approx K_F$. A similar statement holds for the t -channel and u -channel diagrams.

From this point of view, the one-loop contributions are generically power suppressed. The exceptional behavior occurs when K (or Q or Q') is of order Λ/v_F . Unlike the case for large K or Q , the behavior qualitatively differs between the s and t channels.

For the t -channel diagram to make a nonsuppressed contribution to the beta function, the following must hold:

$$\frac{\Lambda}{v_F} < |Q| < \mathcal{O}(1) \times \frac{\Lambda}{v_F}. \quad (17)$$

Thus, there is a window of values where the contribution is nonzero, and the position of the edges of this widow depend on Λ . On the other hand, for the s -channel diagram to make an unsuppressed contribution to the beta function, K must satisfy

$$|K| < \mathcal{O}(1) \times \frac{\Lambda}{v_F}. \quad (18)$$

In particular, $K = 0$ gives an order-one contribution while $Q = 0$ does not. Fig. 4 demonstrates the behavior of the log derivatives assuming a constant coupling.

This difference has a profound effect. In the course of the RG flow, the condition that K or Q is order Λ/K_F changes, since we take Λ to scale down. If K is actually zero from the beginning, there will always be an order-one contribution to the beta function, and this condition is stable throughout the RG flow. This allows attractive couplings with $K = 0$ (the BCS channel) to become strong at small Λ . On the other hand, the condition for the t -channel diagram to give a unsuppressed contribution to the beta function is not stable under the RG flow.




FIG. 4. The log derivatives of the s - and t -channel diagrams measured in units of g^2 for a circular Fermi surface. We assume a constant coupling.

Hence for *any* fixed Q , the t channel only contributes to the beta function for a small period of RG time.

In summary, the contribution to the beta function is power-suppressed throughout the RG flow for generic (large) K and Q . If K or Q is small enough, there are order-one contributions to the beta function, but only for a short RG time. The only exception is “true” BCS scattering, where $|K| < \Lambda/v_F$ throughout the flow. If we assume the UV coupling is weak, this means the only coupling that can be relevant to the ground state instability involves the BCS configuration.

With this context, the startling results for the one-loop VH beta function [Eqs. (9)-(15)] are less surprising. Even with a round Fermi surface, the beta function, and therefore the coupling, depends on K and Q . This is even true for the BCS coupling, which is generically a function of two angular coordinates^{16,18} (playing the role of label

momenta) for noncircular Fermi surfaces.

Finally, the transition from zero contribution to the beta function from the t -channel diagram to a finite contribution as we increase Q from zero is also present for the circular Fermi surface. The major difference in the VH case is the long, flat section of the Fermi surface, which guarantees that the window in Q for which the t channel is not power suppressed is larger than for a circular Fermi surface. Fortunately, we will see that at least for certain observables, we may once again neglect the contribution from the t channel relative to the s channel.

VI. THE LEADING CONTRIBUTION AT ONE LOOP

Section IV demonstrates that at one loop, only the s -channel diagram contains a logarithmic enhancement. Furthermore, the largest possible contribution to the beta function occurs when $K \simeq 0$. This indicates that the kinematic configuration of near-zero net momentum, the BCS channel,²⁶ dominates the low-energy behavior of the theory.

With this in mind, assume the UV the dependence on the external momenta is analytic. This condition will not be preserved under the RG, because the s -channel introduces a nonanalytic dependence on the net momentum K in the four-point coupling. However, if we focus on the BCS configuration we may ignore any nonanalytic dependence on the other momenta to leading-log order.

In the following calculations, we sum the leading VH and NVH contributions. While the precise form of the full results generically depends on the detailed shape of the NVH portion of the Fermi surface, the leading contribution is independent of these details. Instead, this summing procedure turns out to be identical to taking the VH results and replacing the cutoff Υ with V_F , its natural value.

Parameterize the BCS coupling as $g_B(p_1, p_3)$, where p_1 is the label momentum of one of the incoming pair of particles (the other has label momentum $-p_1$) and p_3 is the label momentum of one of the pair of outgoing particles. We find with logarithmic accuracy (see Appendix)

$$\Lambda \frac{dg_B(p_1, p_3)}{d\Lambda} = \frac{1}{4\pi^2} g_B(p_1, 0) g_B(0, p_3) \log \frac{V_F^2}{\Lambda}. \quad (19)$$

An unusual feature of this equation is that the beta function has an explicit dependence on Λ , as well as V_F^2 . The latter can be regarded as an energy scale of order of the bandwidth, $V_F^2 \sim W$. Thus the IR physics retains some information about the UV scale W .

The solution to (19) is

$$g_B(p_1, p_3; \Lambda) = g_B(p_1, p_3; \Lambda_0) - \left(\frac{1}{8\pi^2} \right) \frac{g_B(p_1, 0; \Lambda_0) g_B(0, p_3; \Lambda_0) \left(\log^2 \frac{V_F^2}{\Lambda} - \log^2 \frac{V_F^2}{\Lambda_0} \right)}{1 + \frac{g_B(0, 0; \Lambda_0)}{8\pi^2} \left(\log^2 \frac{V_F^2}{\Lambda} - \log^2 \frac{V_F^2}{\Lambda_0} \right)}. \quad (20)$$

The coupling in the vicinity of the van Hove singularity, $g_B(0,0)$, plays a special role: it “drives” the RG for the other couplings, and when it is attractive at the scale Λ_0 , it sets the one-loop estimate of the strong-coupling scale,

$$\Lambda^* = V_F^2 \exp \left(-\sqrt{\log^2 \frac{V_F^2}{\Lambda_0} + \frac{8\pi^2}{|g_B(0,0;\Lambda_0)|}} \right). \quad (21)$$

As in the ordinary BCS theory²² the strong-coupling scale is non-perturbative in $g(\Lambda_0)$. However, the usual dependence of this scale on the microscopic parameters differs from (21). While (21) simplifies considerably if we set $\Lambda_0 = V_F^2 \sim W$, this choice may be unphysical if the van Hove EFT is obtained by integrating out some other degrees of freedom at a scale below W . For example, if the short-range interaction arises both from the screened Coulomb repulsion and the phonon-mediated attraction, the van Hove EFT applies only up to energy scales of order the Debye frequency ω_D , which is usually much smaller than the bandwidth W . Then the natural choice for Λ_0 is ω_D , and we have a hierarchy of scales $V_F^2 \simeq W \gg \omega_D$.

To understand some of the limitations of this formalism, consider the amplitude (as opposed to the beta function) in the BCS configuration. If we assume a momentum-independent BCS coupling, it is straightforward to evaluate the one-loop amplitude with logarithmic accuracy:

$$\mathcal{A}_{\text{BCS}}(E) = \frac{g_B^2}{8\pi^2} \left(\log^2 \frac{V_F^2}{\Lambda} - \log^2 \frac{V_F^2}{E} - i\pi \log \frac{V_F^2}{E} \right), \quad (22)$$

where we have kept only the leading terms in the real and imaginary parts. Taking the log derivative of equation (22) with respect to Λ reproduces the beta function (19) for $g_B(0,0)$. However, the imaginary part of the amplitude depends on $\log \frac{V_F^2}{E}$. This large log is not resummed by the standard beta function and indicates that something akin to the rapidity renormalization group introduced in²⁰ would be necessary to resum subleading logs.

In the special case $\Upsilon^2 = \Lambda$ our scheme in the VH region resembles that of Ref.⁵. In that work it is implicitly assumed that g is repulsive, and that Λ can be taken as high as the bandwidth, so that the NVH region is effectively absorbed into the VH region. However, lowering Λ then results in integrating some low-energy modes and requires nonlocal counterterms.

VII. HIGHER-ORDER RENORMALIZATION

Let us discuss how higher-order corrections modify Eq. (22). This is particularly important for the kinematic configuration with zero net momentum which controls the Cooper instability. Since the beta function at zero net momentum contains a logarithm of a large ratio,




FIG. 5. The two-loop self-energy with finite imaginary part.

$\log(V_F^2/\Lambda)$, one may wonder if the one-loop computation is reliable in this kinematic configuration, or if one needs to resum the logs in the beta function itself. We will call logs containing V_F^2 , such as $\log(V_F^2/E)$ or $\log(V_F^2/\Lambda)$, rapidity logs. We want to estimate the contribution of higher rapidity logs to the beta function at zero net momentum.

We will limit ourselves to the analysis of 2-loop diagrams. We take $\Upsilon \sim V_F$, in which case there are no large rapidity logs in non-VH loops. The renormalized coupling g is related to the bare coupling g_b by

$$g_b = g Z_4 Z_2^{-2}, \quad (23)$$

where Z_4 is the renormalization factor for the particle-particle four-point amplitude, and Z_2 is the wave function renormalization. Z_2 is finite at one loop, and at two-loop order is determined from the on-shell behavior of the self-energy diagram, Fig. 5, whose imaginary part is finite even without the rapidity cut-off^{1,2}, and therefore does not contain rapidity logs.

Two-loop contributions to Z_4 arise from diagrams such as in Fig. 6. (Iterations of one-loop diagrams do not contribute since their infinities are removed by one-loop counter-terms). Their contributions to the beta function can be estimated using what we already know about the one-loop diagrams. For example, the diagram Fig. 6a is the obtained from the one-loop s -channel diagram by replacing one of the vertices with the one-loop t -channel diagram. The latter does not contain rapidity logs, so the contribution of the whole diagram to the beta function should behave in the same way as that of the one-loop s -channel diagram. In particular, it contains at most a single $\log(V_F^2/\Lambda)$ at zero net momentum. The diagram Fig. 6b can be regarded as a one-loop t -channel diagram with one vertex replaced with a one-loop s -channel diagram. The latter amplitude contains at most two rapidity logs, so the contribution of Fig. 6b to the beta function contains at most $\log^2(V_F^2/\Lambda)$. We conclude that with logarithmic accuracy the two-loop beta function at zero net momentum has the form

$$\beta(g) = \frac{1}{4\pi^2} (g^2 + Cg^3) \log \frac{V_F^2}{\Lambda} + C' g^3 \log^2 \frac{V_F^2}{\Lambda}, \quad (24)$$

where C and C' are constants.

Now we can see if the resummation of rapidity logs in the beta function is necessary. Eq. (20) indicates that the one-loop RG equations resum logs of the form $g \log^2(V_F^2/\Lambda)$. Thus we are assuming that




FIG. 6. Examples of two-loop contributions to the beta function. The diagrams with iterated loops are not shown. The diagram on the right can contribute a double rapidity log to the beta function.

$g \log^2(V_F^2/\Lambda) \lesssim 1$, while $g \log(V_F^2/\Lambda) \ll 1$. This implies $g^3 \log^2(V_F^2/\Lambda)$ is parametrically suppressed relative to $g^2 \log(V_F^2/\Lambda)$. We conjecture that this behavior persists at higher loops, in the sense that every extra power of g is accompanied by at most a single rapidity log. If this is true, then resumming the rapidity logs in the beta function will not change qualitative conclusions regarding the RG flow and the Cooper instability.

VIII. THE COLLINEAR REGION AS A MARGINAL FERMION LIQUID

By definition, the collinear region is the part of the VH region where $|p_+|$ is of order of the rapidity cutoff Υ , while $|p_-|$ is less or equal than Λ/Υ . The anticollinear region is defined similarly, but with p_+ and p_- exchanged. Each of the following statements regarding the collinear region also applies to the anticollinear region.

Everywhere in the collinear region, the Fermi velocity is nonzero. Naively, one might conclude that this region is no different from the NVH region. In particular, one might think that the usual Fermi surface EFT¹⁶ applies both in the NVH and the collinear region, but this is incorrect. To see why, recall that canonical Fermi surface EFT predicts that all interactions (apart from forward and BCS scattering) are irrelevant, and thus the quasi-particle width scales like $E^2/v_F k_B$ for small E . In the NVH region, v_F is of order W/k_B , thus the Fermi liquid theory applies for E much smaller than W . But it is well known¹⁶ that additional marginal interactions arise when a portion of the Fermi surface is related to another portion of the Fermi surface by a translation in momentum space (nesting). The translation vector Q is called the nesting vector. The collinear region is an extreme exam-

ple of this, since the Fermi surface is approximately invariant with respect to arbitrary shifts with $Q = (Q_+, 0)$. Following Wilczek and Nayak²³, we will refer to such a Fermi surface as flat.

Wilczek and Nayak emphasized the failure of the Fermi liquid theory for flat Fermi surfaces and proposed that the correct EFT for flat Fermi surfaces is quasi-1D, with the component of momentum parallel to the Fermi surface playing the role of a continuous label. In particular, the four-fermion interaction is marginal for generic combinations of momenta rather than irrelevant.

But there is also an important difference between the collinear region and the model of interacting 1D fermions (the Luttinger model). In the Luttinger model, the coupling is exactly marginal (has vanishing beta function). This is most easily seen using bosonization, which turns the Luttinger model into a free boson with a linear dispersion law. The vanishing of the beta function does not apply to the EFT describing the collinear region. The reason is that, unlike in the 1D case, the Fermi velocity varies along the Fermi surface. For definiteness, let us consider the collinear region and set $\mu = 0$. Then the “small” component of momentum is p_- , while the “large” one is p_+ . If we treat p_+ as a continuous label, the Fermi velocity is

$$v_F(p_+) = p_+. \quad (25)$$

As long as we consider generic scattering events between particles for which p_+ is $\mathcal{O}(\Upsilon)$, the four-fermion coupling can be Taylor expanded in p_- , but not in p_+ . Thus the leading interaction term,

$$S_{int} = \int dt \int d^2 p_1 d^2 p_2 d^2 p_3 \times \frac{1}{4} g(p_{1+}, p_{2+}, p_{3+}) \epsilon^{\alpha_1 \alpha_2} \epsilon^{\alpha_3 \alpha_4} \psi_{\alpha_1} \psi_{\alpha_2} \psi_{\alpha_3}^\dagger \psi_{\alpha_4}^\dagger, \quad (26)$$

depends on a function of three real variables $g(p_{1+}, p_{2+}, p_{3+})$ which we take to be spin independent. This choice of spin structure for the interaction corresponds to the spin-singlet coupling, which we will focus on here. Furthermore, we take g to be symmetric under $p_1 \leftrightarrow p_2$ and $p_3 \leftrightarrow p_1 + p_2 - p_3$ independently, so the vertex factor is

$$i(\delta_{\alpha_1 \alpha_3} \delta_{\alpha_2 \alpha_4} - \delta_{\alpha_1 \alpha_4} \delta_{\alpha_2 \alpha_3}) g(p_1, p_2, p_3). \quad (27)$$

It is straightforward to compute the beta function for g . We find:

$$\begin{aligned}
\frac{dg(p_{1+}, p_{2+}, p_{3+})}{d \log \mu} = & \frac{1}{2\pi^2} \int_K^\Upsilon dq \frac{g(p_{1+}, p_{2+}, q)g(q, K - q, p_{3+})}{2q - K} \\
& + \frac{1}{8\pi^2 Q} \int_{Q-\min(Q, \Upsilon)}^{\min(Q, \Upsilon)} dq g(p_{1+}, q, p_{3+})g(p_{2+}, q + Q, p_{4+}) \\
& + \frac{1}{8\pi^2 Q'} \int_{Q'-\min(Q', \Upsilon)}^{\min(Q', \Upsilon)} dq g(p_{1+}, q, p_{4+})g(p_{2+}, q + Q', p_{3+}) \quad (28)
\end{aligned}$$

plus terms suppressed by $\varepsilon(p)/\Lambda$, where p is one of the external momenta. Here $K = p_{1+} + p_{2+}$ and $Q = p_{1+} - p_{3+}$ are assumed to be positive, for definiteness. Even if we take g to be independent of the “large” components of momenta at some scale, the RG evolution is nontrivial and introduces momentum dependence. At higher orders we will also have to take into account the renormalization of the Fermi velocity function $v_F(p_+)$. Finally, we neglected the spin-triplet coupling. Even if it is set to zero in the UV, it will be generated by radiative corrections, and thus a renormalizable theory should have both couplings. The above computation which takes into account only the spin-singlet coupling merely illustrates our point that the beta functions are nonzero in the collinear region.

The EFT which includes only the collinear region is sufficient to compute the width of the quasiparticle whose momentum is in the collinear region, where $|p_+|$ is of order Υ . If one formally takes the limit $\Upsilon \rightarrow \infty$ and assumes that the coupling g is independent of momenta, the leading-order computation can be performed in the toy model and gives^{1,2}:

$$\Gamma(E) \sim g^2 E. \quad (29)$$

The linear dependence on E follows from dimensional analysis and is a hallmark of the marginal Fermi liquid¹³. The computation in the toy model cannot be extended to higher orders, since it is not a renormalizable theory. However, if we include the NVH region by introducing the rapidity cutoff $\Upsilon = V_F$, dimensional analysis gives a similar result:

$$\Gamma(E) \sim h(|p_+|/V_F)E, \quad (30)$$

where V_F is the typical Fermi velocity in the NVH region. At leading (two-loop) order the function $h(x)$ is of order g^2 , but is not a constant even if one assumes, for simplicity, that g is a constant. Evaluating the imaginary part of the self-energy diagram (Fig. 5) numerically, we find the result in Fig. 7.

The expression (30) is valid provided we can neglect the chemical potential μ which is a relevant coupling. Thus it holds in the range $|\mu| \ll E \ll W$. The corrections are of several sorts. The NVH region contribution is of order E^2/W , as usual. The corrections from a nonzero




FIG. 7. Numerical results for the dependence of h on $|p_+|/V_F$ in units of g^2 assuming a constant coupling. h is normalized to g^2 for $V_F \rightarrow \infty$.

μ are of order μ^2/E . Finally, higher orders in perturbation theory will give the function h a weak (logarithmic) dependence on E .

One of the defining properties of the MFL is that the quasiparticle width, defined via the imaginary part of the on-shell self-energy, is proportional to energy. The above arguments show that the Marginal Fermi Liquid behavior¹³ is a robust consequence of the proximity to a van Hove singularity. On the other hand, the dependence of the width on the “large” component of momentum can be nontrivial, unlike in the simplest models of Marginal Fermi Liquids.

IX. CONCLUSIONS

We have presented a systematic effective field theory description of systems with a van Hove singularity. The formalism is valid to leading power in an expansion in E/W and generalizes the classic results in^{16,17}. We have shown that the theory is renormalizable with all counterterms being local in the sense that they are finite in the zero energy limit. That such a formalism exists had to be the case given that any well-defined microscopic local theory must yield a renormalizable description, if it is properly formulated. A crucial ingredient in generating

such a theory is the inclusion of all the relevant modes on the Fermi surface. Given that the entire surface is necessarily part of the IR description of the theory, it is not surprising that focusing solely on one region leads to nonlocalities.

The EFT that we constructed depends on a coupling function $g(k_1, k_2, k_3, -k_1 - k_2, -k_3)$ that cannot be expanded in powers of momenta (except when all momenta are ultrasoft). The appearance of an arbitrary function of six variables makes the theory much less predictive than the usual Fermi surface RG which has two marginal couplings which depend on two variables each (for a 2D Fermi liquid). Nevertheless, we showed that in the BCS channel the EFT can be greatly simplified, provided we keep only logarithmically-enhanced terms. In this channel, one is left with a single function of two variables which satisfies a simple RG equation.

We have utilized our formalism to show that generic theories with van Hove singularities will lead to Marginal Fermi Liquid behavior as previously anticipated using toy models^{1,2}. This behavior arises in both the soft and collinear subsectors of the VH region, the latter of which can constitute a considerable fraction of the Fermi surface. Thus our conclusions disagree with⁵, where it was argued that for $E \gg \mu$ the Fermi liquid picture is valid. Our treatment of the collinear region clarifies the physics of Fermi surfaces with flat regions as discussed in²³. We also show that the running of the coupling in the BCS channel is logarithmically enhanced, and the coupling itself runs double logarithmically, in agreement with^{3,4}.

Appendix: One-loop beta function calculations

We consider only the spin-singlet interaction. The interaction part of the Lagrangian is

$$\frac{g}{4} \epsilon^{ab} \epsilon^{cd} \psi_a^\dagger \psi_b^\dagger \psi_c \psi_d. \quad (\text{A.1})$$

The tree-level four-point amplitude is

$$S_{abcd} g, \quad (\text{A.2})$$

where

$$S_{abcd} = \epsilon_{ab} \epsilon_{cd} \quad (\text{A.3})$$

is the spin structure of the amplitude. Momentum conservation means the coupling is a function of three momenta (six real variables) $g(p_1, p_2, p_3)$ which we take to be symmetric under $p_1 \leftrightarrow p_2$ and $p_3 \leftrightarrow p_1 + p_2 - p_3$.

The one-loop four-point amplitude has contributions from the s , t , and u channels. The s channel has the same spin-singlet structure as the tree-level amplitude, but the t and u channels generically contain both spin-singlet and spin-triplet contributions. To simplify this analysis, we ignore the spin-triplet contributions entirely. Then the amplitude takes the form

$$\mathcal{A}_{abcd} = S_{abcd} (\mathcal{A}_S + \mathcal{A}_T + \mathcal{A}_U). \quad (\text{A.4})$$

The u -channel amplitude follows from the t -channel amplitude via exchange of the two outgoing momenta. Besides the momentum dependence of the coupling, the s - and t -channel diagrams only depend on the external momenta through $K = p_1 + p_2$ and $Q = p_1 - p_3$ respectively.

After performing the energy integrals via contours, changing coordinates to $p_\pm = p_x \pm p_y$, and assuming time-reversal invariance for the dispersion,

$$\begin{aligned} \mathcal{A}_S = & -\frac{1}{8\pi^2} \int d^2k \frac{\theta_{S+} - \theta_{S-}}{\varepsilon_k + \varepsilon_{K-k} - E - i\epsilon \text{sign } \varepsilon_k} \\ & \times g(p_1, p_2, k) g(k, K - k, p_3) f(k) f(K - k), \end{aligned} \quad (\text{A.5})$$

where

$$\theta_{S\pm} \equiv \theta(\pm\varepsilon_k) \theta(\pm\varepsilon_{K-k}) \quad (\text{A.6})$$

constrain the momenta of the loop propagators. E is the net energy of the external particles. f contains all information regarding the cutoffs:

$$f(k) = \theta(\Lambda - |\varepsilon_k|) \theta(\Upsilon - |k_+|) \theta(\Upsilon - |k_-|). \quad (\text{A.7})$$

Similarly, the t -channel amplitude is

$$\begin{aligned} \mathcal{A}_T = & \frac{1}{16\pi^2} \int d^2k \frac{\theta_{T+} - \theta_{T-}}{\varepsilon_k - \varepsilon_{k+Q} + E_T - i\epsilon \text{sign } \varepsilon_k} \\ & \times g(p_1, k, p_3) g(k + Q, p_2, k) f(k) f(k + Q), \end{aligned} \quad (\text{A.8})$$

with

$$\theta_{T\pm} \equiv \theta(\pm\varepsilon_k) \theta(\mp\varepsilon_{k+Q}). \quad (\text{A.9})$$

E_T is the transfer energy of the external particles. The extra factor of $1/2$ arises from isolating the spin-singlet contribution.

We are interested in the beta function for g . Taking the logarithmic derivative with respect to Λ yields

$$\Lambda \frac{d\mathcal{A}_S}{d\Lambda} = -\frac{1}{4\pi^2} (I_{S+} + I_{S-}), \quad (\text{A.10})$$

$$\Lambda \frac{d\mathcal{A}_T}{d\Lambda} = \frac{1}{8\pi^2} (I_{T+} + I_{T-}), \quad (\text{A.11})$$

where

$$\begin{aligned} I_{S\pm} \equiv & \pm \Lambda \int d^2k \frac{\delta(\Lambda \mp \varepsilon_k) \theta(\pm\varepsilon_{K-k})}{\varepsilon_k + \varepsilon_{K-k}} \\ & \times \theta(\Upsilon - |k_+|) \theta(\Upsilon - |k_-|) f(k - K) \\ & \times g(p_1, p_2, k) g(k, K - k, p_3), \end{aligned} \quad (\text{A.12})$$

$$\begin{aligned} I_{T\pm} \equiv & \pm \frac{1}{2} \Lambda \int d^2k \frac{\delta(\Lambda \mp \varepsilon_k) \theta(\mp\varepsilon_{k+Q})}{\varepsilon_k - \varepsilon_{k+Q}} \\ & \times \theta(\Upsilon - |k_+|) \theta(\Upsilon - |k_-|) f(k + Q) \\ & \times [g(p_1, -k - Q, p_3) g(-k, p_2, -k - Q) \\ & + g(p_1, k, p_3) g(k + Q, p_2, k)]. \end{aligned} \quad (\text{A.13})$$

We have dropped E and E_T because they lead to power-suppressed terms in the beta function. The remaining integrals are similar to each other. They involve integrating over the one-dimensional space where one of the particles in the loop has $\varepsilon = \pm\Lambda$ and the other has either the same sign for ε (for the s channel) or the opposite sign (for the t channel).

Define P to be equal to K for the s -channel diagram and $-Q$ for the t -channel diagram. We exploit the $O(1,1)$ invariance of the dispersion to replace P in by $\tilde{P} = \sqrt{|\varepsilon_P|}(\text{sign } P_+, \text{sign } P_-)$ in each of the integrals by changing variables:

$$k_+ = k'_+/\eta, \quad (\text{A.14})$$

$$k_- = \eta k'_-, \quad (\text{A.15})$$

with

$$\eta \equiv \frac{\sqrt{|\varepsilon_P|}}{|P_+|} = \sqrt{\left|\frac{P_-}{P_+}\right|}. \quad (\text{A.16})$$

We may take η to be less than one by exchanging k_+ and k_- if necessary. The step functions involving the rapidity cutoff Υ are not invariant under this change of variables. In particular,

$$\theta(\Upsilon - |k_+|) \rightarrow \theta(\eta\Upsilon - |k_+|), \quad (\text{A.17})$$

$$\theta(\Upsilon - |k_-|) \rightarrow \theta(\Upsilon/\eta - |k_-|), \quad (\text{A.18})$$

and similarly for the step functions constraining $k - P$. These set the limits of integration on the remaining k_+ integrals (once we have performed the k_- integrals with the delta function) if the energy constraints do not set stricter limits. With that in mind, let us first analyze the limits in the absence of a rapidity cutoff.

1. Integration limits

We can write generic expressions for the various possible integration limits in each of the four remaining integrals. As before, take P to be either K or $-Q$. Define

$$s_{\pm} = \text{sign } P_{\pm}, \quad (\text{A.19})$$

$$s_k = \begin{cases} 1 & \text{for } I_{S+}, I_{T+}, \\ -1 & \text{for } I_{S-}, I_{T-}, \end{cases} \quad (\text{A.20})$$

$$s_p = \begin{cases} 1 & \text{for } I_{S+}, I_{T-}, \\ -1 & \text{for } I_{S-}, I_{T+}. \end{cases} \quad (\text{A.21})$$

The remaining k_+ integrals have limits at

$$\lambda_A \equiv s_+ \sqrt{|\varepsilon_P|}, \quad (\text{A.22})$$

$$\lambda_B \equiv s_- s_k \frac{\Lambda}{\sqrt{|\varepsilon_P|}}. \quad (\text{A.23})$$

Whenever the quantities

$$\lambda_{\pm} \equiv \frac{1}{2} \sqrt{|\varepsilon_P|} \left(s_+ + s_- (s_k - s_p) \frac{\Lambda}{|\varepsilon_P|} \pm \sqrt{1 - 2s_+ s_- (s_k + s_p) \frac{\Lambda}{|\varepsilon_P|} + (s_k - s_p)^2 \frac{\Lambda^2}{\varepsilon_P^2}} \right) \quad (\text{A.24})$$

are purely real, the remaining integrals also have limits at λ_{\pm} . In that case the integration region splits into two disjoint pieces.

For I_{S+} and I_{S-} , $s_k = s_p$ and Eq. (A.24) simplifies to

$$\lambda_{\pm} = \frac{1}{2} \sqrt{|\varepsilon_P|} \left(s_+ \pm \sqrt{1 - 4s_+ s_- s_k \frac{\Lambda}{|\varepsilon_P|}} \right). \quad (\text{A.25})$$

For I_{T+} and I_{T-} , $s_p = -s_k$ and this simplifies to

$$\lambda_{\pm} = \frac{1}{2} \sqrt{|\varepsilon_P|} \left(s_+ + \frac{2s_- s_k \Lambda}{|\varepsilon_P|} \pm \sqrt{1 + \frac{4\Lambda^2}{\varepsilon_P^2}} \right). \quad (\text{A.26})$$

Therefore the t -channel integrals always split into two pieces. The s -channel integrals split unless

$$|\varepsilon_K| \leq 4\Lambda, \quad (\text{A.27})$$

in which case $I_{S\text{sign } \varepsilon_K}$ is over a single contiguous region bounded by λ_A and λ_B . In that case,

$$I_{S\text{sign } \varepsilon_K} = \int_{\min(\lambda_A, \lambda_B)}^{\max(\lambda_A, \lambda_B)} \dots \quad (\text{A.28})$$

We will see that this integral (and only this one) is generally divergent as $\varepsilon_K \rightarrow 0$, and that this divergence is cured by the rapidity cutoff.

For the remaining three integrals (and for $I_{S\text{sign } \varepsilon_K}$ if (A.27) is not satisfied), the integration regions are bounded on one side by either λ_A or λ_B and on the other by either λ_+ or λ_- . The remaining integrals take the form

$$I = \left(\int_{\lambda_1}^{\lambda_2} + \int_{\lambda_3}^{\lambda_4} \right) dk_+ \dots \quad (\text{A.29})$$

where λ_1 through λ_4 are the limits sorted in ascending order.

2. Rapidity limits

At this point, let us simplify the discussion by taking g to be a momentum-independent constant. We will find that this assumption is not consistent, because the beta function depends on the momentum. Suppressing

the integration limits, the remaining integrals take the following form:

$$I = g^2 \int d^2k \frac{\delta(\varepsilon_k - s_k \Lambda) \theta(s_p \varepsilon_{k-P}) F(k, P)}{1 + s_p \frac{\varepsilon_{k-P}}{\Lambda}}, \quad (\text{A.30})$$

where

$$F(k, P) = \theta(\Upsilon - |k_+|) \theta(\Upsilon - |k_-|) f(k - P). \quad (\text{A.31})$$

Note that the step function constrains the value of the denominator to be between 1 and 2 throughout the integration region, so all of the integrals are nonnegative.

Under $P_+ \rightarrow -P_+$ or $P_- \rightarrow -P_-$, $I_{S+} \leftrightarrow I_{S-}$ and $I_{T+} \leftrightarrow I_{T-}$. Thus we can take both components of K and Q to be positive without loss of generality. This simplification would have held earlier if we had assumed the coupling function obeys particle-hole symmetry, but this symmetry is generically broken by the NVH region.

We can now find the effect of the rapidity cutoff on the integration limits for the various integrals. The lower limit on k_+ imposed by the rapidity cutoff for I_{S+} and I_{S-} is

$$\lambda_{R1} = \eta \max \left(\frac{\Lambda}{\Upsilon}, P_+ - \Upsilon \right), \quad (\text{A.32})$$

where $\eta = \sqrt{|P_-/P_+|}$. The upper limit is

$$\lambda_{R2} = \begin{cases} \eta \Upsilon, & P_- \leq \Upsilon \\ \eta \min \left(\Upsilon, \frac{\Lambda}{P_- - \Upsilon} \right), & P_- > \Upsilon. \end{cases} \quad (\text{A.33})$$

$\lambda_{R1}/\lambda_{R2}$ replaces the lower/upper limits in (A.28) or (A.29) when it is within either integration region. Alternatively, if it is less than/greater than both limits in one of the integrals, the integral is set to zero.

For I_{S-} and I_{T-} , one of the integration regions has negative k_+ and the other has positive k_+ . There are four possible rapidity limits:

$$\lambda_{R3} = \eta(P_+ - \Upsilon), \quad (\text{A.34})$$

$$\lambda_{R4} = -\frac{\eta \Lambda}{\Upsilon}, \quad (\text{A.35})$$

$$\lambda_{R5} = \frac{\eta \Lambda}{\Upsilon - P_-}, \quad (\text{A.36})$$

$$\lambda_{R6} = \eta \Upsilon. \quad (\text{A.37})$$

$\lambda_{R3}/\lambda_{R4}$ replace lower/upper limit for the negative integration region and $\lambda_{R5}/\lambda_{R6}$ replace the lower/upper limit for the positive region, or they set the appropriate integrals to zero, acting in a manner analogous to that described above for λ_{R1} and λ_{R2} .

3. Indefinite integrals

Evaluating the delta function in (A.30) and changing variables to $x = \frac{k_+}{\sqrt{\Lambda}}$ yields

$$I = g^2 \int \frac{dx}{|x|} \frac{1}{1 + s_k - s_p \alpha (x + s_k/x - \alpha)}, \quad (\text{A.38})$$

where

$$\alpha \equiv \sqrt{\frac{|\varepsilon_P|}{\Lambda}}. \quad (\text{A.39})$$

We can directly compute the indefinite integrals for $I_{S\pm}$ and $I_{T\pm}$ as long as we make use of the restrictions on the integration limits implied by Section 1. We find

$$I_{S+}(x) = \frac{g^2}{\sqrt{\alpha^4 + 4}} \times \log \left| \frac{-2 - \alpha^2 + \sqrt{\alpha^4 + 4} + 2\alpha x}{2 + \alpha^2 + \sqrt{\alpha^4 + 4} - 2\alpha x} \right|, \quad (\text{A.40})$$

$$I_{S-}(x) = \frac{g^2 \text{sign } x}{\sqrt{\alpha^4 + 4}} \times \log \left| \frac{-2 + \alpha^2 + \sqrt{\alpha^4 + 4} - 2\alpha x}{2 - \alpha^2 + \sqrt{\alpha^4 + 4} + 2\alpha x} \right|. \quad (\text{A.41})$$

For I_{T+} , the appropriate indefinite integral depends on the magnitude of α , or in other words on the relative size of $|\varepsilon_Q|$ and Λ :

$$I_{T+}(x) = -\frac{2g^2}{\alpha \sqrt{4 - \alpha^2}} \arctan \left(\frac{\alpha - 2x}{\sqrt{4 - \alpha^2}} \right) \quad (\text{A.42})$$

for $|\varepsilon_Q| < 4\Lambda$ and

$$I_{T+}(x) = \frac{g^2}{\alpha \sqrt{\alpha^2 - 4}} \log \left| \frac{\alpha + \sqrt{\alpha^2 - 4} - 2x}{-\alpha + \sqrt{\alpha^2 - 4} + 2x} \right| \quad (\text{A.43})$$

for $|\varepsilon_Q| > 4\Lambda$. Finally,

$$I_{T-}(x) = \frac{g^2 \text{sign } x}{\alpha \sqrt{\alpha^2 + 4}} \log \left| \frac{-\alpha + \sqrt{\alpha^2 - 4} + 2x}{\alpha + \sqrt{\alpha^2 - 4} - 2x} \right|. \quad (\text{A.44})$$

4. Collinear-anticollinear limit

Consider scattering between generic collinear and anticollinear particles. In this case, $|\varepsilon_P| \gg 4\Lambda$, so $\alpha \gg 2$. In the $\alpha \rightarrow \infty$ limit,

$$\frac{1}{\alpha \sqrt{\alpha^2 + 4}} \approx \frac{1}{\alpha \sqrt{\alpha^2 - 4}} \approx \frac{1}{\sqrt{\alpha^4 + 4}} \rightarrow \frac{\Lambda}{|\varepsilon_P|}. \quad (\text{A.45})$$

This suggests the beta function is suppressed by $\Lambda/|\varepsilon_P|$ for collinear-anticollinear scattering, although we must also check the behavior of the log functions.

When $a > 2$, there are always two disjoint integration regions. Exchanging k_+ and k_- exchanges the two regions, so when we ignore the rapidity cutoff they must have the same value. Evaluating the integral between

$$\frac{\lambda_+}{\sqrt{\Lambda}} = \frac{1}{2}\alpha \left(1 + (s_k - s_p) \frac{1}{\alpha^2} + \sqrt{1 - 2(s_k + s_p) \frac{1}{\alpha^2} + (s_k - s_p)^2 \frac{1}{\alpha^4}} \right) \quad (\text{A.46})$$

and

$$\frac{\lambda_A}{\sqrt{\Lambda}} = \alpha, \quad (\text{A.47})$$

reversing the order if $\lambda_+ > \lambda_A$, and taking the $\alpha \rightarrow \infty$ limit yields the same result for each integral:

$$I = (2 \log 2) \frac{\Lambda}{|\varepsilon_P|} g^2 + \mathcal{O} \left(\frac{\Lambda^2}{\varepsilon_P^2} \right) \quad (\text{A.48})$$

for collinear-anticollinear scattering. All one-loop contributions to the beta function are therefore power suppressed in this limit. This remains true when we include the rapidity cutoff, since it can only reduce the size of the integration region. Furthermore, such interactions continue to be power suppressed after we drop the assumption of a momentum-independent coupling, since the integration region always shrinks to zero size as ε_P becomes large.

5. Collinear limit

Consider the scenario where all scattered particles are restricted to the collinear region and assume the external momenta p_{i+} and their sums/differences (K_+ , Q_+ , and Q'_+) are all order Υ and much larger than $\sqrt{\Lambda}$. Furthermore, assume the scattered particles have energies well below the cutoff, so $\varepsilon_p \ll \Lambda$. Together, these imply that the perpendicular components of momenta are small:

$$p_{i-} = \frac{\varepsilon_{p_i}}{p_{i+}} = \mathcal{O} \left(\frac{\varepsilon_{p_i}}{\Upsilon} \right) \ll \frac{\Lambda}{\Upsilon}. \quad (\text{A.49})$$

The following results also hold, with appropriate modifications, if all momenta lie in the anticollinear region.

In this limit, only the rapidity cutoff on the collinear components of momenta comes into play. Furthermore,

$$\varepsilon_K = \varepsilon_{p_1} \left(1 + \frac{p_{2+}}{p_{1+}} \right) + \varepsilon_{p_2} \left(1 + \frac{p_{1+}}{p_{2+}} \right) \ll \Lambda \quad (\text{A.50})$$

since we have assumed that the collinear components of the incoming particles are all of the same order. As a result,

$$\alpha = \sqrt{\frac{\varepsilon_K}{\Lambda}} \ll 1. \quad (\text{A.51})$$

Similar statements hold for the t - and u -channel contributions.

After a change of variables, the integration limits for both the I_{S+} and I_{S-} integrals are $\sqrt{|\varepsilon_K/\Lambda|} = \alpha$ and

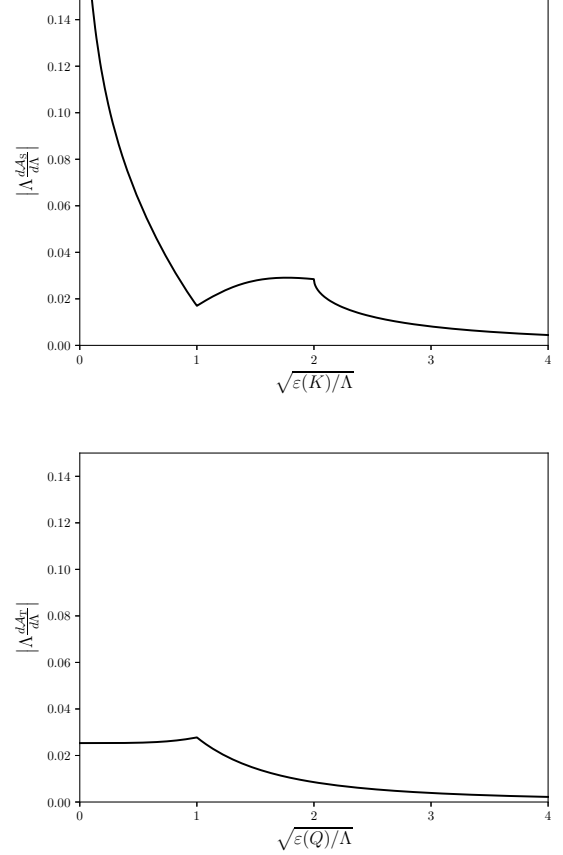


FIG. 8. The log derivatives of the s - and t -channel diagrams in units of g^2 with no rapidity cutoff for a Fermi surface with a van Hove singularity. We assume a constant coupling.

$\eta\Upsilon/\sqrt{\Lambda} = \alpha\Upsilon/|K_+|$. Only one of the integration regions for I_{S-} remains after we impose the rapidity cutoff. Substituting these limits into the indefinite integrals and taking the small- α limit yields

$$I_{S+} \approx I_{S-} \rightarrow \frac{g^2}{2} \log \left(\frac{2\Upsilon}{K_+} - 1 \right). \quad (\text{A.52})$$

For the I_{T+} and I_{T-} integrals, only one of the two integration regions remains. The limits are

$$\frac{\lambda_{\mp}}{\sqrt{\Lambda}} = \frac{1}{2}\alpha \left(1 \pm \frac{2}{\alpha^2} \mp \sqrt{1 + \frac{4}{\alpha^4}} \right), \quad (\text{A.53})$$

with λ_- for I_{T+} and λ_+ for I_{T-} , and $\lambda_A/\sqrt{\Lambda} = \alpha$. Substituting these into the appropriate indefinite integrals

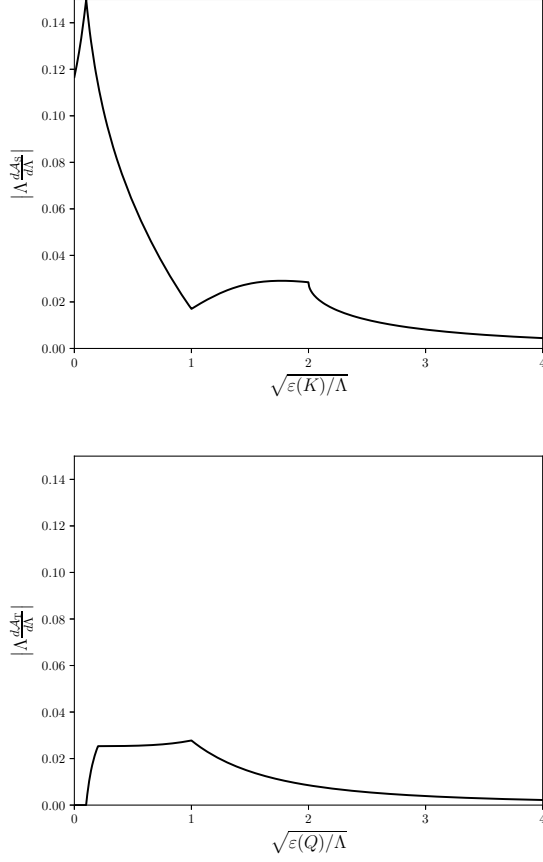


FIG. 9. Plots demonstrating how the rapidity cutoff modifies Fig. (8).

gives

$$I_{T+} \approx I_{T-} \rightarrow \frac{g^2}{2} \quad (\text{A.54})$$

in the collinear limit.

There are two important features of (A.52) and (A.54). First, the contributions from the s , t , and u channels will all be order g^2 . Second, the integrals are independent of the small (anticollinear) components of the external momenta. These conclusions do not depend of our assumption of a momentum-independent coupling. Backtracking through our derivation and restoring the momentum dependence yields (28).

6. Forward scattering

If both components of Q are smaller than Λ/Υ , both integration regions for I_{T+} and I_{T-} shrink to zero size. Thus, the t -channel contribution to the beta function disappears in the forward-scattering limit in the presence of a rapidity cutoff; see Fig. 9. This is analogous to the situation discussed in¹⁶, where the forward scattering function makes no contribution to the beta functions for a

round Fermi surface. As in the case of a round Fermi surface, there is a sharp change in the contribution to the beta function once ε_Q exceeds a threshold; compare Fig. 4 and Fig. 9.

7. BCS limit

Consider I_{S+} in the $K_{\pm} \rightarrow 0$ limit. Since $\varepsilon_K < 4\Lambda$, there is a single contiguous integration region, bounded by $\lambda_A = \sqrt{\varepsilon_K}$ and $\lambda_B = \sqrt{\Lambda/\varepsilon_K}$. The extent of this region diverges as we lower ε_K . We find

$$I_{S+} = \frac{g^2}{\sqrt{\alpha^4 + 4}} \times \log \left(\frac{(-2 + \alpha^2 + \sqrt{\alpha^4 + 4})(\alpha^2 + \sqrt{\alpha^4 + 4})}{(-\alpha^2 + \sqrt{\alpha^4 + 4})(2 - \alpha^2 + \sqrt{\alpha^4 + 4})} \right). \quad (\text{A.55})$$

Taking the small α limit yields

$$I_{S+} \rightarrow \frac{1}{2} g^2 \log \frac{4\Lambda}{\varepsilon_K}, \quad (\text{A.56})$$

which diverges at $\varepsilon_K = 0$. This is the divergence that forced us to introduce the rapidity regulator. I_{S-} has the same value as I_{S+} in the small α limit.

Introducing the rapidity cutoff regulates the divergence. The rapidity cutoff restricts the I_{S+} integral to run from $\eta\sqrt{\Lambda}/\Upsilon$ to $\eta\Upsilon/\sqrt{\Lambda}$ and the I_{S-} integral to run from η to $\eta\Upsilon/\sqrt{\Lambda}$. Plugging these into (A.40) and (A.41) and taking the $\alpha \rightarrow 0$ limit yields

$$I_{\pm} \rightarrow \frac{1}{2} g^2 \log \frac{\Upsilon^2}{\Lambda}, \quad (\text{A.57})$$

so

$$\Lambda \frac{d\mathcal{A}_S}{d\Lambda} = -\frac{g}{4\pi^2} \log \frac{\Upsilon^2}{\Lambda} \quad (\text{A.58})$$

for back-to-back interactions.

8. Generic BCS beta function

The previous results indicate that we may take the coupling for fixed ultrasoft net momentum (the BCS configuration) to be analytic in the other momenta to leading-log order. Furthermore, we may drop all but the s -channel diagram to this order. Parameterize the BCS coupling $g_B(p_1, p_3)$ in terms of one of the incoming momenta p_1 and one of the outgoing momenta p_3 at fixed ultrasoft K . The log derivative of the amplitude is (A.10), with

$$I_{S\pm} \equiv \pm \Lambda \int d^2k \frac{\delta(\Lambda \mp \varepsilon_k) \theta(\pm \varepsilon_{K-k}) F(k, K)}{2\varepsilon_k - E} \times g_B(p_1, k) g_B(k, p_3) \quad (\text{A.59})$$

and $F(k, K)$ from (A.31). Take the components of K to be positive but infinitesimal to avoid ambiguity from the




FIG. 10. The integration regions for $I_{S\pm}$.

definition of the step functions. Eq. (A.59) receives contributions from several one-dimensional regions of momentum space; see Fig. 10.

Call I_{++} the contribution from the region with $k_+ > k_- > 0$. Evaluating the k_- integral with the delta function yields

$$I_{++} = \frac{1}{2} \int_{\sqrt{\Lambda}}^{\Upsilon} \frac{dk}{k} g_B(p_1, k) g_B(k, p_3) \quad (\text{A.60})$$

up to power-suppressed terms. Assume it is possible to expand the coupling function in k_+ and k_- . The resulting expression for I_{++} will include terms of the form

$$\int_{\sqrt{\Lambda}}^{\Upsilon} \frac{dk_+}{k_+} k_+^m \left(\frac{\Lambda}{k_+} \right)^n \partial_{k_+}^m \partial_{k_-}^n [g_B(p_1, 0) g_B(0, p_3)]. \quad (\text{A.61})$$

The natural scale for the derivatives is $(1/V_F)^{m+n}$. As a result, terms with $m \neq n$ give at most order-one contributions to the beta function. When $m = n$, terms

in (A.61) take the form

$$\left(\frac{\Lambda}{V_F^2} \right)^n \int_{\sqrt{\Lambda}}^{\Upsilon} \frac{dk_+}{k_+} = \frac{1}{2} \left(\frac{\Lambda}{V_F^2} \right)^n \log \frac{\Upsilon^2}{\Lambda}. \quad (\text{A.62})$$

Since we assume $\Lambda \ll V_F^2$, these are suppressed unless $n = 0$. The $n = m = 0$ term is log enhanced, and the leading-log result is therefore

$$I_{++} = \frac{1}{4} g_B(p_1, 0) g_B(0, p_3) \log \frac{\Upsilon^2}{\Lambda}. \quad (\text{A.63})$$

A similar analysis holds for each of the the terms in $I_{S+} + I_{S-}$. Adding the NVH region cancels the Υ dependence. Finally, setting the log derivative with respect to Λ of the sum of the tree-level amplitude $g_B(p_1, p_3)$ and the one-loop amplitude equal to zero implies

$$\Lambda \frac{dg_B(p_1, p_3)}{d\Lambda} = \frac{1}{4\pi^2} g_B(p_1, 0) g_B(0, p_3) \log \frac{V_F^2}{\Lambda}. \quad (\text{A.64})$$

The beta function for the the coupling between modes in the vicinity of the VH point, $g_V \equiv g_B(0, 0)$, is independent of the other couplings, and the solution is

$$g_V(\Lambda) = \frac{g_V(\Lambda_0)}{1 + \frac{g_V(\Lambda_0)}{8\pi^2} \left(\log^2 \frac{V_F^2}{\Lambda} - \log^2 \frac{V_F^2}{\Lambda_0} \right)}. \quad (\text{A.65})$$

Using this, the beta function for $g_B(p_1, 0)$ becomes

$$\begin{aligned} \Lambda \frac{dg_B(p_1, 0)}{d\Lambda} &= \frac{1}{4\pi^2} \frac{g_B(p_1, 0) g_V(\Lambda_0) \frac{V_F^2}{\Lambda}}{1 + \frac{g_V(\Lambda_0)}{8\pi^2} \left(\log^2 \frac{V_F^2}{\Lambda} - \log^2 \frac{V_F^2}{\Lambda_0} \right)}, \end{aligned} \quad (\text{A.66})$$

with solution

$$g_B(p_1, 0; \Lambda) = \frac{g_B(p_1, 0; \Lambda_0)}{1 + \frac{g_V(\Lambda_0)}{8\pi^2} \left(\log^2 \frac{V_F^2}{\Lambda} - \log^2 \frac{V_F^2}{\Lambda_0} \right)}. \quad (\text{A.67})$$

An analogous result holds for $g_B(0, p_3; \Lambda)$. Substituting these into the beta function for $g_B(p_1, p_3)$ and solving yields

$$g_B(p_1, p_3; \Lambda) = g_B(p_1, p_3; \Lambda_0) - \left(\frac{1}{8\pi^2} \right) \frac{g_B(p_1, 0; \Lambda_0) g_B(0, p_3; \Lambda_0) \left(\log^2 \frac{V_F^2}{\Lambda} - \log^2 \frac{V_F^2}{\Lambda_0} \right)}{1 + \frac{g_V(0, 0; \Lambda_0)}{8\pi^2} \left(\log^2 \frac{V_F^2}{\Lambda} - \log^2 \frac{V_F^2}{\Lambda_0} \right)}. \quad (\text{A.68})$$

We see that the expressions for $g_B(0, 0; \Lambda)$, $g_B(p_1, 0; \Lambda)$, and $g_B(0, p_3; \Lambda)$ are in fact special cases of this general result.

ACKNOWLEDGMENTS

This work was supported by the DOE contracts DE-SC0011632, DE-FG02-92ER40701, DOE-ER-40682-143, and DE-AC02-6CH03000. The authors gratefully ac-

knowledge helpful conversations with Joe Polchinski, Max Metlitskii, Lesik Motrunich, and Ingmar Saberi. AK and TM are grateful to the Simons Center for Geometry and Physics for hospitality during various stages of this work. AK is also grateful to the Aspen Center for

Physics, the Kavli Institute for Physics and Mathematics of the Universe, and Institut des Hautes Etudes Scientifiques for hospitality. IZR is grateful to the Caltech theory group for hospitality and to the Moore Foundation for support.

* kapustin@theory.caltech.edu

† tmckinney@caltech.edu

‡ izr@andrew.cmu.edu

¹ S. Gopalan, O. Gunnarsson and O. K. Andersen, Phys. Rev. B **46**, 11798 (1992).

² P. C. Pattnaik, C. L. Kane, D. M. Newns and C. C. Tsuei, Phys. Rev. B **45**, 5714 (1992).

³ I. Dzialoshinskii, J. Phys. France **6**, 119 (1996).

⁴ J. V. Alvarez, J. Gonzalez, F. Guinea, and M. A. H. Vozmediano, J. Phys. Soc. Jpn. **67**, 1868 (1998).

⁵ D. Menashe and B. Laikhtman, Phys. Rev. B **59**, 13592 (1999).

⁶ C. Honerkamp, M. Salmhofer, N. Furukawa, and T. M. Rice, Phys. Rev. B **63**, 035109 (2001).

⁷ D. Zanchi and H. J. Schulz, Phys. Rev. B **61**, 13609 (2000) [arXiv:cond-mat/9812303].

⁸ C. J. Halboth and W. Metzner, Phys. Rev. B **61**, 7364 (2000) [arXiv:cond-mat/9908471].

⁹ V. Yu. Irkhin, A. A. Katanin, and M. I. Katsnelson, Phys. Rev. B **64**, 165107 (2001).

¹⁰ G. Kastinakis, Physica C **340**, 119 (2000).

¹¹ G. Kastinakis, Phys. Rev. B **71**, 014520 (2005).

¹² R. Nandkishore, L. Levitov, and A. Chubukov, Nature Physics **8**, 158 (2012).

¹³ C. M. Varma, P. B. Littlewood, S. Schmitt-Rink, E. Abrahams and A. E. Ruckenstein, Phys. Rev. Lett. **63**, 1996 (1989).

¹⁴ For a review of the van Hove scenario see R. S. Markiewicz, J. Phys. Chem. Solids **58**, pp. 1179-1310 (1997); J. Bok and J. Bouvier, in *New Topics in Superconductivity Research*, ed. B. P. Martin, Nova Science Publishers (2006).

¹⁵ G. Benfatto and G. Gallavotti, Phys. Rev. B **42**, 9967

(1990).

¹⁶ R. Shankar, Physica A **177**, 530 (1991), Rev. Mod. Phys. **66**, 129 (1994).

¹⁷ J. Polchinski, in *Boulder 1992, Proceedings, Recent directions in particle theory* 235-274 [arXiv: hep-th/9210046].

¹⁸ S. Weinberg, Nucl. Phys. B **413**, 567 (1994) [cond-mat/9306055].

¹⁹ S. Ghamari, S.-S. Lee, and C. Kallin, Phys. Rev. B **92**, 085112 (2015).

²⁰ J. Y. Chiu, A. Jain, D. Neill and I. Z. Rothstein, Phys. Rev. Lett. **108**, 151601 (2012) [arXiv:1104.0881 [hep-ph]]. J. -Y. Chiu, A. Jain, D. Neill and I. Z. Rothstein, JHEP **1205**, 084 (2012) [arXiv:1202.0814 [hep-ph]].

²¹ M. E. Luke, A. V. Manohar and I. Z. Rothstein, Phys. Rev. D **61**, 074025 (2000) [hep-ph/9910209].

²² R. Schrieffer, *Theory of superconductivity*, Perseus Books, revised edition (1999).

²³ C. Nayak and F. Wilczek, arXiv:cond-mat/9408016.

²⁴ Note that Υ breaks the $O(1,1)$ symmetry but preserves the particle-hole symmetry.

²⁵ Note that for the s-channel diagram, taking $\varepsilon_K \gtrsim \Lambda$ is equivalent to injecting a large virtuality into the loop, which is formally outside the range of validity of the effective theory. The effects of such modes in intermediate states are properly accounted for in higher dimensional, power-suppressed, operators. This is consistent with the result in (15).

²⁶ In this context, the term “BCS” means back-to-back up to an ultrasoft momentum. Generic configurations of only ultrasoft modes therefore qualify as BCS.